Fragmented Justification

Inaugural-Dissertation

zur Erlangung des Doktorgrades der Philosophie der

Ludwig-Maximilians-Universität München

Vorgelegt von Hannah Pillin
aus Kyoto, Japan

2026

Referent: Prof. Dr. Dr. Hannes Leitgeb
Korreferent: Prof. Dr. Christian List
Datum der mündlichen Prüfung: 09.02.2026

This dissertation is dedicated to women and all minorities and marginalised groups. In Philosophy, and everywhere. Especially in times like these, we shall never lose hope, keep believing and, most importantly, continue to take action towards a more equal and just future.

Preface

This dissertation, as well as the path towards it, would not have been possible without many. I would like to take the opportunity to acknowledge and thank those people.

I would like to express my deep and heartfelt gratitude to my primary supervisor, Hannes Leitgeb, Chair of Logic and Philosophy of Language, and Co-Director of the Munich Center for Mathematical Philosophy, LMU Munich, and secondary supervisor Christian List, Chair of Philosophy and Decision Theory, and Co-Director of the Munich Center for Mathematical Philosophy, LMU Munich.

I look up to their way of doing and thinking about philosophy; both the clarity of thought, philosophical insight and creativity, as well as the benevolence with which they engage in philosophical debates with scholars and students equally, never cease to inspire me. I chose to pursue my PhD at the MCMP because of this, and I could not have asked for better supervisors, both academically, as well as personally, or a more collaborative, friendly and fruitful environment than the MCMP for my doctoral studies.

Moreover, I would like Liam Kofi Bright, Associate Professor at the Department of Philosophy, Logic and Scientific Method at the London School of Economics and Political Science, who was my thesis supervisor during my masters studies LSE. I am especially grateful for Liam's mentorship beyond the graduation of my masters and throughout my doctoral studies.

Next, I would like to thank Ursula Danninger, administrator at the Munich Center for Mathematical Philosophy at LMU Munich. Ursula's support throughout my time at LMU is invaluable to me. I would like

to especially thank her for the very nice working relationship within the "Mentoringprogramm" of the philosophy faculty at LMU.

I would like to thank Norbert Gratzl, Akademischer Oberrat at the MCMP, LMU, and Paolo Mancosu, Willis S. and Marion Slusser Professor of Philosophy at the University of California at Berkeley for helping to make it possible for me to spend the spring semester of 2019 as a visiting student at UC Berkeley. I would also like to thank Norbert for helping to make it possible for me to pursue a masters degree at the London School of Economics at the same time as pursuing a second masters degree at the MCMP, LMU. I would also like to thank Dirk Kindermann, Adjunct Lecturer at the University of Vienna, who kindly hosted me during a research stay at the University of Vienna in the fall of 2023. Furthermore, I would like to thank Agustín Rayo, Professor of Philosophy, Kenan Sahin Dean of MIT School of Humanities, Arts, and Social Sciences, at Massachusetts Institute of Technology, for hosting my stay at MIT during the 2024 fall semester.

I would like to thank audiences at various conferences, workshops and colloquiums. I thank audiences at the weekly logic colloquium, work in progress seminar, decision theory colloquium, as well as graduate seminar mathematical philosophy, all at the MCMP, for helpful questions and comments. I also thank audiences at the 2023 Knowledge in Crisis Colloquium at University of Vienna, the 2024 Graduate Conference in Philosophy at UC Irvine, the LOFT 15, 15th Conference on Logic and the Foundations of Game and Decision Theory in Bayreuth, and the weekly WIP colloquium, as well as graduate student colloquium at MIT.

Furthermore, I would like to thank several anonymous reviewers of article versions of chapters 3 and 4 of this dissertation. At this point, I would like to note that, as of August 2025, an article version of chapter 4 of this dissertation has been published with Synthese – An International Journal for Epistemology, Methodology and Philosophy of Science, Springer Publishing House (Pillin, 2025[1]). Moreover, an article version of chapter 3

[1]The contents of the paper and the corresponding dissertation chapter overlap significantly, the main difference lying in the exposition of the content, one has the form

is, at time submission to UB München, after defense, has been accepted (forthcoming) for publication by Philosophical Studies – An International Journal for Philosophy in the Analytic Tradition.

I would like to thank all my peers, friends and colleagues at the MCMP, as well as the Philosophy PhD Student Stammtisch Group. Thank you for the philosophical discussions, debates, lunches, coffee breaks, and social activities; thank you for your community and camaraderie. You essentially contributed to the great time I had here at the MCMP!

I would also like to thank the students I had the honour of having in my tutorials on the introduction to logic, and introduction to epistemology and philosophy of science. Teaching helped me see philosophy through a different lens, and it was a great pleasure to think about philosophical questions together with my students.

I would also like to say a special thank you to the women and members of marginalised groups I was and am able to look up to within academic philosophy. Being a woman in philosophy is not always easy for me, always being in the minority in a seminar room, or a conference, and finding the courage to continuously speak up and share my thoughts nevertheless, can be very challenging. Being able to see, identify with, and connect with other women in philosophy is very meaningful to me.

Moreover, thank you to Dr. Margit Dirscherl, Akademische Rätin at Institut für Deutsche Philologie, LMU Munich, who helped me to receive my first opportunity to give a talk at an academic conference. During my bachelors degree, in 2018, she encouraged me to submit an essay on a piece of poetry by Heinrich Heine to the 21st Forum Junge Heine Forschung at the University of Düsseldorf that got accepted and became my first academic conference talk.

Thank you, furthermore, to Inge Schmidt, my mathematics teacher in high school, for recommending me for the Deutsche Schülerakademie – Science Academy Baden-Württemberg – in 2011, where I first encountered logic and analytic philosophy. And thank you for recommending me for the

of a dissertation chapter, the other one has the form of a journal article.

Studienstiftung des Deutschen Volkes.

Thank you to my violin teacher, Charlotte Balle, who was a big influence for me throughout my childhood and adolescence.

Before coming to the more personal acknowledgements, I would like to thank the Studienstiftung des Deutschen Volkes. They have supported my studies since I started my bachelors degree and have helped me pursue my studies in a way that would not have been possible in the same way otherwise. I also thank them for the PhD scholarship I have received throughout my doctoral studies which made it possible for me to focus on my research without having to worry about finances.

Then, on a more personal level, thank you for the friends I have been able to make through the Studienstiftung.

I would like to thank my parents, Gita Finkenbeiner and Mathias Pillin. Thank you for gifting me life, raising me the way you did, motivating me to follow and explore my passions, and supporting me in going after my dreams. Thank you both from the bottom of my heart for your support and love. And Mama, thank you for the enormous effort it must have been for you to raise us the absolutely exceptional way you did as a single mom. Thank you to my siblings Josef Pillin and Zoe Pillin, I am grateful to be the big sister of two such beautiful human beings. Furthermore, thank you to my entire family and extended family (with a special mention to my two deceased grandmothers, Maria and Herta) from the bottom of my heart. Thank you for your support and love throughout my entire life. Danke! I also want to thank my childhood pets, Fee, Neko, and Anton – I miss you.

I want to say a heartfelt thank you to all my amazing friends, with special thanks to Maria Vasilaki and Arina Chiru. Thank you for your love and friendship. I would also like to thank Klara Fuhr, Leonora Erhardt, Patricia Schöllhorn-Gaar, Julian Mattlinger, Sarah Kramer, Pauline Becker, Carlotta Nardi, Eileen Andert, my knitting group of friends, Johanna, Sophia, Alysha, and the Brussels Rue Jonniaux family, Edeline, Pietro, Lukas, Ana. Thank you for your friendship. I also thank my current orchestra and all orchestras throughout the years, as well as the community at Caritas Münchner Korbinian Suppenküche.

Lastly, I would like to thank myself for following my interests, curiosity and intuition, believing in myself, working towards my goals, and for standing up after difficulties and disappointments to keep going forward, try again and do better.

And finally, especially since part of this dissertation will be on the preface paradox, let me say a few philosophical introductory words: I am fallible and believe that there will be an error in this dissertation. At the same time, I justifiedly believe that each claim I make throughout this dissertation is true. These two statements will not conflict, because I am fragmented with respect to justification. Throughout this dissertation, it will become clear what this means.

Abstract

This dissertation develops a novel account of fragmented epistemic[2] justification. It departs from and expands on the literature on fragmented accounts of belief, and develops a theory where not belief, but the justification of belief, is fragmented. This account is novel, and an addition to the already existing literature on fragmented accounts of belief. It will also shed light on and help to develop answers to some underdiscussed and unsolved questions in the existing literature on fragmentation, such as questions about the rationality of fragmentation.

In a nutshell, what it means for justification to be fragmented will be that some piece of evidence might justify a certain proposition relative to some mental fragment, whereas that same piece of evidence might not justify that same proposition relative to a different mental fragment. The epistemic and non-epistemic circumstances, such as differences in accessible information, or different perspectives, focus, or practical goals etc., of the agent can vary across fragments, which is why these circumstances will be the parameters distinguishing fragments from each other. What exactly this means will become clear throughout this dissertation.

For the first qualitative part of the dissertation, justification will be modelled to be a non-monotonic justification relation that holds between pieces of evidence and the propositions they justify, which will be indexed to fragments. Fragments, here, will be modelled by using total plausibility pre-orders over possible worlds.

Later on, within a quantitative account of fragmented justification,

[2]Unless explicitly noted otherwise, from here on, if I use the word "justification", I always refer to epistemic justification.

there will be a variety of ways to spell out the concept of justification. Here, fragments will be modelled by using (subjective) probability functions or credence functions.

After an introductory chapter, there will be a chapter focusing on an overview over relevant literature that is being discussed or worked with in this dissertation. Then, chapter 3 presents a skeleton of a qualitative (non-probabilistic) account of fragmented justification and shows how this account is able to provide a novel and attractive solution to the preface paradox (Makinson, 1965), one of the central paradoxes of epistemology.

Then, chapter 4 develops a more full fledged account of qualitative fragmented justification, using this time not the preface paradox but a scenario from climate science as a motivating example for the account. Here, a scientific expert about climate science seems to be justified in believing an overall inconsistent set of propositions. The chapter explains how the toolkit of orthodox epistemology is limited in dealing with situations like this, and argues that fragmentation is a more appropriate strategy than existing approaches to deal with (seemingly rational) inconsistent doxastic states. The chapter provides a formalisation of this example, as well as a number of justifying reasons for why fragmentation can be rational. It, also, develops some precautionary norms of fragmentation. It concludes by explaining how fragmented justification relates to belief, and argues that fragmented justification provides a novel rationality argument for accounts of fragmented belief, and, furthermore, how one should go about de-fragmenting in light of practical rationality.

Chapters 5 and 6 turn to a probabilistic or quantitative account of fragmented justification – fragmented Bayesianism. Chapter 5 presents a very general version of fragmented Bayesianism, much more general than the existing accounts of fragmented Bayesianism that already exist. The basic idea is that the doxastic state of an agent will be represented not by one credence function, but by a set of credence functions, corresponding to different fragments. The focal point of this chapter is a list of parameters, such as background knowledge, algebra, confirmation measure, or ur-prior, that are part of the orthodox Bayesianism machinery, which, each, can be

chosen to be either fragmented or global within a fragmented account of Bayesianism. Depending on which and how many of those parameters are chosen to be kept global or to render fragmented, the resulting account of fragmented Bayesianism will be more or less severely fragmented. The chapter provides motivations and arguments for why Bayesianism in general, and each of the parameters in particular, deserve fragmented versions. It also provides a brief comparison between the qualitative account of fragmented justification developed in the previous chapters with fragmented Bayesianism.

Lastly, chapter 6 compares fragmented Bayesianism with another framework, imprecise Bayesianism. Imprecise Bayesianism aims to model agents who do not have enough information to form precise credences. Fragmentation, on the other hand, aims to represent agents with fragmented doxastic states, where fragments are individuated via differences in the agent's epistemic and non-epistemic circumstances. The formal models underlying these two accounts are strikingly similar. In both frameworks, the agent's doxastic state is represented not via one but via many different credence functions. The literature on fragmentation and imprecision do not comment on the striking similarity between these two different accounts. This is why chapter 6 is devoted to presenting and comparing their similarities and differences.

The dissertation closes with a conclusion and some remarks on potential avenues for further work.

Deutsche Zusammenfassung

Diese Dissertation entwickelt eine neuartige Theorie fragmentierter Rechtfertigung. Sie weicht von der bestehenden Literatur über fragmentierte Überzeugungen ab und erweitert diese, indem sie eine Theorie entwickelt, in der nicht die Überzeugungen (beliefs), sondern die Rechtfertigung dieser Überzeugungen eine fragmentierte Struktur hat. Dieser Ansatz ist neuartig und stellt eine Ergänzung zur und Erweiterung der bereits bestehenden Literatur über Theorien von fragmentierten epistemischen Attitüden, wie zu Beispiel fragmentierter Überzeugung, dar. Zudem werden in dieser einige bislang wenig diskutierte und ungelöste Fragen der Fragmentationsliteratur besprochen und beantwortet, insbesondere Fragen in Bezug auf die Rationalität von Fragmentierung.

Fragmentierte Rechtfertigung erlaubt es, dass bestimmte Daten / Evidenz eine Proposition relativ zu einem bestimmten mentalen Fragment rechtfertigen können, während dieselben Daten dieselbe Proposition relativ zu einem anderen Fragment möglichwerweise nicht rechtfertigen. Mentale Fragmente unterscheiden sich durch unterschiedliche epistemischen und nicht-epistemische Umstände, in denen sich ein Agent relativ zu den unterschiedlichen Fragmenten befinden kann, wie, unter anderem, die Information, die zugämglich ist, die Perspektive des Agenten, oder praktische Ziele, die verfolgt werden. Was genau das bedeutet, wird im Laufe der Dissertation klar werden. Ich werde auf bestimmte Fragment-individuierende Parameter genauer eingehen, und erklären, warum solch Fragment-individuierende Parmeter Fragmentierung auch rational rechtfertigen können. Diese Theorie wird es möglich machen, bestimmte Probeme zu lösen, mit denen traditio-

nelle Erkenntnistheorie Probleme hat. Der Grund hierfür ist, dass die Theorie, die ich in dieser Dissertation entwickle, es schaffen wird, zu erklären, warum bestimmte epistmeische Situationen rational sein können, obwohl traditionelle Erkenntnistheorie solche Situationen als irrational kennzeichnen würden.

Im ersten qualitativen Teil dieser Dissertation wird Rechtfertigung als nicht-monotone Rechtfertigungsrelation zwischen Daten / Evidenz (repr. durch Propositionen) und den Propositionen, die zu rechtfertigen sind, modelliert. Diese Rechtfertigungsrelation ist Fragment-relativ. Das bedeutet, dass die Validität einer bestimmten Instanz dieser Relation von der Beschaffenheit des Fragments abhängt, auf dem die Funktion definiert ist. Fragmente werden hier durch totale Plausibilitätsordnungen über mögliche Welten modelliert. Das bedeutet, in Kurzform, dass

Später, im hinteren Teil der Dissertation, biete ich, zusätzlich zu der qualitativen Theorie von fragmentierter Rechtfertigung auch eine quantitative / probabilistische Version davon an. Der Grund hierfür ist, dass es, zusätzlich zu qualitativen nicht-framentierten Ansätzen von Rechtertigung in der Literatur bisher, auch sehr einflussreiche quantitative / proabilistische Ansätze von Rechtfertigung gibt, am bekanntesten der sogennante Bayesianismus.

Im Rahmen einer quantitativen Theorie fragmentierter Rechtfertigung, kann das Konzept der Rechtfertigung auf verschiedene Weise ausgearbeitet werden, worauf ich in den späteren Kapiteln dieser Dissertation genauer eingehen werde. In diesem probabilistischen Teil werden Fragmente mithilfe von subjektiven Wahrscheinlichkeitsfunktionen modelliert werden.

Nach einem einleitenden Kapitel, in dem ich meine persönliche, als auch die philosophische Motivation für das Thema dieser Dissertation motiviere, folgt ein Überblickskapitel über die relevante Literatur, die in der Dissertation behandelt und verwendet wird. Das Kapitel rekapituliert zunächst in Form einer kurzen Zusammenfassung das historische Interesse in das Konzept der Rechtfertigung, und gibt anschließend einen Überblick über die gängigen Theorien der Rechtfertigung in der erkenntnistheoretischen Literatur. Das Kapitel gibt, außerdem, eine Einführung in “fragmentier-

te Erkenntnistheorie", da sich diese Dissertation in diesen Literaturstrand eingliedern wird, und ein Verständis von und Überblick über dieselbe essenziell für den Inhalt dieser Dissertation ist. Außerdem folgen im selbigen Kapitel ein Hintergrund zu und Überblick über die Sozialwahltheorie, Bayesianismus, sowie die Theorie der impräzisen Wahrscheinlichkeiten, da alle diese Ansätze im Verlauf der Dissertation relevant werden. Am Ende dieses Kapitels steht eine Zusammenfassung und ein Ausblick darüber, was diese Dissertation für Additionen und Neuerungen zu den jeweiligen Literaturfeldern beitragen wird.

Als nächstes präsentiert Kapitel 3 dann das Grundgerüst einer qualitativen (nicht-probabilistischen) Theorie fragmentierter Rechtfertigung und zeigt, wie dieser Ansatz eine neuartige und überzeugende Lösung für das Vorwort-Paradoxon (Makinson, 1965) bietet – eines der zentralen Paradoxien der Erkenntnistheorie.

In dem Vorwort-Paradoxon geht es darum, dass ein*e scheinbar hochrationaler Autor*in eines Buches zu einem bestimmten Thema in diesem Buch eine große Menge an gut gerechtfertigen Aussagen $A_1, ..., A_n$ kommuniziert, aber dann im Vorwort desselbigen Buches zugibt, ebenfalls auf Basis guter Rechtfertigung, dass eine der Aussagen $A_1, ..., A_n$ falsch ist. Das scheint irrational, denn damit wird eine global inkonsistente Menge von Aussagen kommuniziert. Der Grund, warum diese Situation nun als Paradox gehandelt wird, ist, dass der oder die Autor*in sehr wohl rational scheint. Alle Aussagen, die in dem Buch getätigt werden (einschließlich des Vorwortes), sind gut gerechtfertigt. Aber traditionelle Erkenntnisthorie besagt, dass eine solch inkonsistente Menge von Aussagen nicht rational sein kann.

Die Theorie der fragmentierten Rechtfertigung als Lösung zu diesem Paradox setzt auf anderen Lösungen auf, die ebenfalls Fragmentierung als Werkzeug verwenden. Die Lösung, die in dieser Dissertation vorgestellt wird, nutzt jedoch die neuartige Theorie der fragmentierten Rechtfertigung, und wird sich aus unterschiedlichen Gründen als besonders attraktive Lösung herausstellen.

Kapitel 4 entwickelt diese qualitative Theorie fragmentierter Rechtferti-

gung dann weiter, diesmal nicht anhand des Vorwort-Paradoxons, sondern mithilfe eines Beispiels aus der Klimawissenschaft. Hierbei scheint ein*e wissenschaftliche*r Expert*in im Gebiet der Klimawissenschaft gerechtfertigt zu sein, in eine insgesamt inkonsistente Menge von Überzeugungen zu haben. Das Kapitel erklärt, warum das Instrumentarium der traditionellen Erkenntnistheorie nicht ausreicht, um solche Situationen richtig zu analysieren und zu lösen, und argumentiert, dass fragmentierte Rechtfertigung eine geeignetere Strategie darstellt als bestehende Ansätze, um mit (scheinbar rationalen) inkonsistenten doxastischen mentalen Zuständen umzugehen. Es liefert eine Formalisierung des Beispiels sowie eine Liste von Gründen, warum Fragmentierung rational sein kann. Zudem werden Vorsichtsnormen im Bezug auf Fragmentierung entwickelt; sie erklären wann und wie Agenten ihr doxastisches System fragmentieren dürfen. Das Kapitel schließt mit einer Erläuterung, wie fragmentierte Rechtfertigung sich zu dem Konzept der Überzeugung verhält. Es wird hier argumentiert, dass fragmentierte Rechtfertigung ein neues Rationalitätsargument für fragmentierte Theorien der Überzeugung darstellt, und analysiert, wie ein Prozess der De-Fragmentiertung aussehen kann, falls praktische Rationalität einen solchen nötig machen sollte.

Kapitel 5 und 6 wenden sich einer probabilistischen bzw. quantitativen Theorie fragmentierter Rechtfertigung zu – dem fragmentierten Bayesianismus.

Kapitel 5 präsentiert eine sehr allgemeine Version des fragmentierten Bayesianismus, die deutlich allgemeiner ist als bestehende Ansätze. Die Grundidee ist, dass der doxastische Zustand eines Agenten nicht durch eine einzige, sondern durch mehrere Wahrscheinlichkeitsfunktionen dargestellt wird – je eine für jedes mentale Fragment. Der Schwerpunkt dieses Kapitels liegt auf einer Liste von Parametern (z.B. Hintergrundwissen, Algebra, Bestätigungsmaß, "Ur-Prior"), die in der orthodoxen bayesianischen Theorie festgelegt sind, und in einer fragmentierten Version jeweils global (unfragmentiert) oder fragmentiert werden können. Je nachdem, welche und wie viele dieser Parameter global festgelegt, oder fragmentiert werden, ergibt sich eine unterschiedlich stark fragmetierte Version einer Theorie

des fragmentierten Bayesianismus. Das Kapitel liefert Argumente, warum Bayesianismus allgemein, und insbesondere jeder dieser Parameter, eine fragmentierte Version verdient. Außerdem wird ein kurzer Vergleich zwischen der zuvor entwickelten qualitativen Theorie fragmentierter Rechtfertigung und der des fragmentierten Bayesianismus gezogen.

Schließlich wird in Kapitel 6 der fragmentierten Bayesianismus mit einer anderen, aber verwandten, Theorie – dem impräzisen Bayesianismus (imprecise Bayesianism) verglichen. Dieser zielt darauf ab, Agenten zu modellieren, die nicht genügend Informationen besitzen, um präzise probabilistische Überzueugungen (Credences) zu bilden. Fragmentierung hingegen modelliert Agenten mit fragmentierten doxastischen Zuständen, wobei sich die Fragmente durch Unterschiede in epistemischen und nicht-epistemischen Umständen des Agenten voneinander unterscheiden. Obwohl die formalen Modelle in beiden Theorien sehr ähnlich sind – in beiden Fällen wird das doxastische System von Agenten durch eine Menge von Wahrscheinlichkeitsfunktionen dargestellt – wird diese Ähnlichkeit in der Literatur nicht thematisiert. Daher widmet sich Kapitel 6 einem Vergleich beider Ansätze.

Die Dissertation schließt mit einem Fazit und einigen Bemerkungen zu möglichen Richtungen für zukünftige Forschung zum Thema der fragmentierten Rechtfertigung ab.

Contents

Chapter 1

Introduction

What distinguishes a rational belief from an irrational one? What is the difference between a rational belief and knowledge? Does the answer to these questions depend on context, that is, does it matter whether we are dealing with beliefs in or knowledge about scientific hypotheses, or, for example, mundane everyday beliefs about the weather?

These are examples of questions that have accompanied me throughout my studies of philosophy, and have even motivated to pursue a degree in philosophy in the first place. Throughout highschool, I would always be interested in how we can know the theories we learn are true, and how we can justify our belief in a hypothesis or theory.

During 8th grade of my high school studies, I got the opportunity to take part in a summer school for mathematically interested students; and I ended up taking a two-week long summer course in logic, titled "Alles logisch? Mit Mathematik und Philosophie dem Denken auf der Spur". I learned about logic and how mathematics and philosophy are so closely related, although one might think differently at first glance. We learned about epistemology and philosophy of language, just to name some topics, and started to understand how logic and mathematics relates to these philosophical disciplines and theories therein. In the end, we reached the overall goal of the course which was to prove the soundness and completeness of propositional logic.

In the first stages of my undergraduate degree in philosophy, I focused on studying logic and philosophy of language. With the years passing, I also developed an increasing interest in the philosophy and methodology of science. Here, the topics that interested me the most could be classified – unbeknownst to me at the time – as epistemology of science: How does the problem of induction pose difficulties to a deductive approach about whether evidence can or cannot validate or support scientific hypotheses? Or, how have different theories of confirmation looked liked through the history of philosophy of science?

I realised – one of the central concepts of interest to philosophers of science, but also to myself, is justification. And as I continued my studies, I learned more about the concept of justification, now also within orthodox epistemology. And I got increasingly interested in what philosophy can tell us — as humans — about how we should reason, how we should reason to be rational human beings, how we should reason to be able to pursue the ultimate goal of getting closer to the truth.

I find these questions especially fascinating, first, because they can have a real impact for my own, and potentially (and hopefully) many other peoples lives. But, moreover, I find them fascinating, because in these questions, so many areas of philosophy come together; it becomes necessary to employ and combine many different sub-disciplines of philosophy to try and tackle questions like these, logic, probability theory, orthodox epistemology, philosophy of science, decision theory, social epistemology, and potentially many more.

I wrote my undergraduate thesis on a topic in the philosophy of language, on epistemic modals, and seemingly paradoxical sentences of the form "It is raining and it might not be". I chose this topic, because I liked how it combined a variety of topics I was interested in. To think about this topic, I needed logic, and semantics in particular; I used epistemic modal logic and dynamic variants of epistemic modal logic to investigate these types of sentences. I also needed to build on literature in the philosophy of language. But this was not all; epistemology was relevant too, after all, the types of modals I was thinking about were epistemic. And the result of this

thesis foreshadows the kind of philosophy I have been interested in subsequently and has, eventually, led me to write this doctoral dissertation. The result of my undergraduate thesis was that the orthodox "propositionalist" (standard possible world semantics) way of modelling sentences including epistemic modals was too strong or demanding and ultimately not suited to appropriately model the semantics and pragmatics of epistemic modals.

Other dimensions that came into the picture during my graduate studies in philosophy were, first, topics in the philosophy of logic, with the main driving questions for me being: Does classical logic prescribe me how to reason in order to be a rational human being? If yes, what if I fail to adhere to the norms that logic prescribes? But if I fail to perfectly adhere to those norms, am I automatically irrational? Is everybody, even really smart scientific experts, irrational, because really, no human can manage to adhere to the demands classical logic poses? Because no one's beliefs are logically closed. And no one's beliefs are always consistent.

Another dimension I became increasingly interested in is how practical and theoretical rationality relate; I started studying decision theory and social choice theory, social epistemology, and dealt with epistemological questions that arise at the intersection between science and policy making and practical decision making more generally.

And centrally, again, I was interested in the concept and role of justification, in relation to all of these topics mentioned. I wrote a first masters dissertation, titled "Can a group be justified in believing a contradiction?". This dissertation brought together my interest in logic, in particular epistemic logic, philosophy of science, orthodox epistemology and social epistemology. I defended the idea that groups can in fact be justified in believing contradictions which is a thesis that does not align with the norms orthodox epistemology imposes for both individual epistemic agents and group agents.

I wrote a second masters dissertation on logics of evidence-based beliefs that reflects my interest in what makes beliefs of epistemic agents rational or justified, and how this can be formally represented and studied.

Fast forward to now writing the introduction of my doctoral dissertation.

A dissertation in epistemology that departs from orthodox epistemology, by adding to the literature on fragmented epistemology, and challenges aspects orthodox epistemology. A dissertation that employs elements from the toolbox of logic (especially non-monotonic logics), as well as probability theory. A dissertation developing a novel theory of justification; a theory of the concept that accompanied me throughout my studies of philosophy. A dissertation that deals with the question of what rationality can look like for for real agents that are fallible and imperfect, and do not adhere to the norms of orthodox epistemology, but still seem rational. A dissertation that thinks about rationality, in both epistemic and practical terms. And a dissertation that combines elements from orthodox epistemology with elements from contemporary formal epistemology.

Now you might be left thinking – ok, but what is this dissertation about more precisely? Let me start at the beginning.

Imagine an author of a book about baby nutrition. This author is a scientific expert about baby nutrition. They have done extensive research on the topic. In this book, they communicate their beliefs about baby nutrition, in a number of chapters that each contain many claims about baby nutrition. The author believes these claims about baby nutrition to be true. They do so because these claims are very well justified by evidence. At the same time, in the preface of their book, this author communicates their belief that one of the claims about baby nutrition they make in the main part of the book is false. This is a variant of the well-known preface paradox, first formulated by Makinson (1965). Taken together, all the claims that the author makes within the book are inconsistent. In the main part of the book, the author communicates a number of claims, $A_1, ..., A_n$ and in the preface of the book, they claim the negation of the conjunction of these propositions $\neg(A_1, ..., A_n)$ which means nothing other than that they claim that one of the claims they make in the main part of the book, $A_1, ..., A_n$, is false.

Orthodox epistemology tells us that if an agent has an overall inconsistent set of beliefs, as the author seems to be doing, then they are irrational. The reason why this is considered a paradox is that intuitively, it seems

that the author is nonetheless rational. Each of the claims they make in the book are well-justified, which, traditionally, is an indicator of the rationality of a belief. And the author is a scientific expert about the field they are writing about, so who could be more rational when it comes to baby nutrition than them?

Fragmented epistemology provides a solution to this problem. Fragmented epistemology departs from the norms of orthodox epistemology in the following way: Orthodox epistemology thinks of the belief state of an agent as a set of beliefs that has to satisfy certain criteria in order for the agent to qualify as rational. For the agent to quality as rational, their set of beliefs has to be consistent (i.e. may not contain claims that contradict each other), and closed under logic / under consequence (i.e. if an agent believes a claim, and another claim logically follows from that claim, they ought to also (explicitly) believe that second claim). The author of the book about baby nutrition does not adhere to these norms, since their overall belief set is inconsistent. Therefore, according to orthodox epistemology, the author is not rational.

Now, according to fragmented accounts of belief, the norms the norms of epistemic orthodoxy do not have to hold for the agent's overall belief set. Instead, fragmented accounts of belief allow for the agent's belief set to be compartmentalised into different fragments, different subsets of the overall set of beliefs that are independently accessible, may get activated in different circumstances, and may contain different kinds of information. In particular, pieces information that are incompatible with each other may be stored in different fragments without making the agent necessarily irrational.

Fragmented accounts of belief can solve the preface paradox by rendering belief a fragment-relative notion. The claims the author make about baby nutrition are contained in / hold relative to one mental fragment, whereas the claim the author makes in the preface about there being a false claim in the book is contained in / holds relative to a different mental fragment. Because the claims that, taken together, produce an inconsistent belief set, are contained in different fragments, this inconsistency can be prevented.

The reason for this is that, for fragmented accounts of belief, orthodox rationality norms hold within but not across fragments. This means, in particular, that closure under conjunction does not hold across fragments, which prevents the author from explicitly believing the outright inconsistent conjunction $A_1 \wedge ... \wedge A_n \wedge \neg(A_1 \wedge ... \wedge A_n)$.

Although fragmented accounts of belief provide a solution to the preface paradox, there are some open issues. Can fragmented belief accounts explain why it would be rational for the author to be fragmented in this way? Makinson motivates the intuition for the author being rational in believing an overall inconsistent set of claims via the good justification that each of the claims enjoy. The justification of the author's beliefs, however, do not play a significant enough role in the solution that fragmented belief accounts provide to the paradox.

The account of fragmented justification I develop in this dissertation builds on the literature already existent on fragmented belief. The account of fragmented justification will be able to provide a novel solution to this paradox that is able to answer the questions just mentioned that fragmented belief accounts do not manage to answer. This will be subject of chapter 3 of this dissertation.

Similar situations, where agents seem to be rational in entertaining an overall inconsistent belief set, also arise in other circumstances beyond the preface paradox. For example, as I develop in chapter 4 of this dissertation, it can happen to scientific experts about complex systems such as the climate when they engage in forming beliefs based on the results of computer modelling simulations. There are situations where the toolkit of orthodox epistemology will be ill-suited to account for the intricacies of specific epistemic situations, like the preface paradox, or a scientific expert justifiedly believing two conflicting hypotheses. I will show how, in such a case, it does not make sense to suspend judgment, nor "forcefully" remove one of the hypotheses from the doxastic state of the agent, which are the options orthodox epistemology provide. I will explain how fragmentation will be much better suited to accommodate for situations like this.

I will explain what factors can distinguish fragments from each other,

as well as under what circumstances fragmentation can be rational. This is a very central aspect of this dissertation; fragmented justification, and fragmentation more generally, should not be seen as an account of rationalising mistakes. I will show how fragmentation will, in some cases, be the epistemically superior strategy compared to the toolkit of orthodox epistemology, which is a novel argument in the fragmentation literature. I will also develop a set of precautionary norms on how and when fragmentation can be rationally permissible.

Beyond, the account of fragmented justification will add to the literature of, both justification, and fragmentation, because all existing orthodox accounts of justification are implicitly based on an unfragmented picture of belief and rationality more generally[1]. I also discuss how fragmented justification relates to belief, and how, given fragmented justification, one can de-fragment in such a way to result in unfragmented beliefs. This will be important because, as I will show, fragmentation can (but does not necessarily) pose significant problems for practical rationality.

The first part of the dissertation – chapters 3 and 4 – will be based on a qualitative (non-probabilistic) conception of justification. Here, justification will be modelled via non-monotonic justification relations that hold between pieces of evidence and the claims they support. These relations will be fragment-relative. Fragments, in this framework, will be modelled to be total plausibility pre-orders over possible worlds. This will result in it being possible that a certain piece of evidence, E, can justify a claim, A, relative to some mental fragment, while E might not justify A relative to another mental fragment.

The second part of the dissertation – chapters 5 and 6 – will be devoted to probabilistic notions of justification and fragmented versions thereof. Here, different mental fragments will be modelled by using different subjective probability functions or credence functions. In chapter 5, I develop a very general account of fragmented Bayesianism. Moreover, in chapter 6, I focus on another phenomenon that is closely related to, but distinct from,

[1] I give an argument for this in section 2.7.1.

fragmentation: Imprecision. I will compare fragmented Bayesianism with models of imprecise probabilities. The reason for why this is important and relevant to this dissertation is that the formal details of both frameworks are almost identical, at least strikingly similar. And the interpretations of the formalisms exhibit interesting parallels and differences which are deserving of a discussion that has not yet taken place in the literature.

At this point, one might ask – ok, justification is an important concept. But why fragmentation? Why is the development of a fragmented account of justification the topic of this dissertation?

I am fascinated by fragmented epistemology, because it makes it possible to formulate models of epistemic attitudes that provide rationality norms that are relaxed in comparison to orthodox epistemology while still being demanding enough for those models to remain normatively significant, instead of merely improving on descriptive accuracy. I find this fascinating, since it allows us to think about a concept of justification and a notion of rationality that is still normatively demanding, but remains "attainable" for real agents. Measured by the norms of orthodox epistemology, all human agents are inevitably irrational. Contrary to this stands the intuition that despite being fallible and making mistakes, there are real agents that do reason in a rational way, even if their reasoning fails to adhere to the rationality criteria orthodox epistemology imposes. Fragmentation is able to capture this intuition, and fragmented justification will be able to shed more light on it.

Lastly, before giving an overview over the structure of this dissertation, let me explain why I am primarily interested in justification (and the justification of belief), and not a different epistemic attitude, such as, for example, knowledge. Firstly, I subscribe to the so-called reasons-first paradigm[2], as opposed to the Williamsonian knowledge-first paradigm (T. Williamson, 2000). For the former, the primitive and foundational notion in epistemology on which all other more complex notions build is the notion of a reason, or a justification. Here, more complex notions that build on the notion of

[2]For an overview over the merits or this epistemological paradigm, cf. eg. Schmidt, 2020.

a reason are the notion of belief, and even more demanding and complex, the notion of knowledge. For the latter, the Williamsonian knowledge-first paradigm, it is the opposite. Here, knowledge is the primitive epistemological notion and the other notions, such as belief and justified belief are defined using knowledge as a primitive.

For reasons-first epistemology, knowledge is the most demanding concept, in the sense that it has the most and most demanding epistemic conditions that need to be satisfied in order for an agent to qualify as knowing something. There are many different notions of (reasons-first) knowledge proposed and discussed in the epistemological literature. Besides all of them being very demanding for real agents to satisfy, they have another thing in common. The vast majority of definitions of knowledge agree that a necessary condition of knowledge is justified (true) belief, and then they diverge in what else, beyond that, is needed for knowledge. Now, often, in practice, both in everyday life and in science, justified belief (and not necessarily knowledge) is already a really useful and fruitful concept, both in epistemic terms, but also concerning practical rationality. I do not, of course, deny that it is interesting and fruitful to try to define the concept of knowledge. But justified belief, and justification in particular, is both a necessary condition of many definitions of knowledge, but also useful in its own right, independent of the notion of knowledge – which is why I focus on this concept. The following chapter will shed more light on why I think this is the case.

This dissertation will be structured as follows:

Chapter 2 provides an overview over and background to the literature that is relevant for this dissertation, including justification, orthodox models of belief, fragmented epistemology, social choice and judgment aggregation, Bayesianism and imprecise probabilities. At the end of chapter 2, there will, furthermore, be a detailed overview over the kinds of contributions this dissertation makes to the literature in each of these fields.

Chapter 3 presents a new solution to the preface paradox (Makinson, 1965). It makes use of a fragmented account of justification, and thus, unlike existing solutions, employs the central concept involved in the paradox

– justification. First, I will present the original paradox, and recall a prominent family of solutions to it – solutions that make use of fragmentation accounts of belief. I will argue that, while fragmented belief solutions go in the right direction, what does the work in this paradox is not belief but justification. I then present my account of fragmented justification. I show how it solves the preface paradox in a novel way and I provide reasons for why it is a good solution. The upshot of my account of fragmented justification will be that the justification relations that hold between pieces of evidence and the propositions they justify are fragment-relative.

Chapter 4 further develops the account of fragmented account of justification. The account is motivated via an example different from the preface paradox – this time via an example from climate science. Within the IPCC assessment report for climate change, an overall inconsistent set of seemingly justified beliefs seems to be asserted. The main reason why this happens will be the extreme uncertainty involved in learning about complex systems such as the climate. The chapter explains how the toolkit of orthodox epistemology is limited in and not suited for dealing with situations like this, and argues that fragmentation is a more appropriate strategy than existing approaches to deal with (seemingly rational) inconsistent doxastic states. After presenting the example, the chapter gives a formalisation of this example. The upshot of this formalisation is that justifications are represented via non-monotonic fragment-indexed justification relations that hold between propositions that need justification, and evidence. The chapter explains why fragmentation is rationally warranted here, and develops an explicit set of precautionary norms that should guide the fragmentation process.

It, furthermore, argues that the account of fragmented justification developed here is compatible with fragmented belief, and how the account of fragmented justification can serve as a novel argument for the rationality of fragmented belief accounts developed in the literature.

And lastly, the chapter explains how sometimes, even though fragmentation might be epistemically rational, de-fragmentation will become necessary in light of practical rationality. It shows how de-fragmentation can

work for epistemically rational fragmented agents, which is a novel discussion within the literature on fragmentation.

Chapter 5 moves from a qualitative account of fragmented justification to a probabilistic account thereof. It presents a very general version of a fragmented account of Bayesianism as an epistemic theory of justification. First, it recalls the central theses of orthodox Bayesianism, and provides an overview over the notions of justification that can be formulated within orthodox Bayesianism. It presents a very general account of fragmented Bayesianism, and then provides a list of parameters along which the account can be made moderately or very fragmented, depending on which and how many of those parameters one chooses to relativise to fragments. It, furthermore, presents an argument for fragmented Bayesianism. I first provide a very general argument for fragmentation, and the provide arguments in favour of the fragmentation of each of the parameters. Lastly, the chapter concludes by a brief comparison between the account of fragmented Bayesianism developed here and the qualitative account of fragmented justification developed earlier.

Chapter 6 deals with how fragmented Bayesianism relates to the account of imprecise Bayesianism. Imprecise Bayesianism, as opposed to orthodox Bayesianism, aims to model agents who do not have enough information to form precise credences. Under this account, an agent's doxastic state is allowed to be made up of not one but many different credence functions, to model imprecision in the assignment of credences to propositions. Fragmentation, on the other hand, aims to represent agents with fragmented doxastic states, where fragments are individuated via differences in the agent's epistemic and non-epistemic circumstances, such as differences in accessible information, or different practical goals and decisions. The formal model underlying this account is very similar to imprecise Bayesianism: The agent's doxastic state is represented via not one but many different credence functions.

The literature on fragmentation and imprecision do not comment on the striking similarity between these two different accounts. This chapter is devoted to presenting and comparing their similarities and differences.

The upshot will be that the main difference lies in, first, the interpretation of the formalisms, how these two accounts deal with questions of practical rationality, as well as how they deal with new information entering the doxastic state and updating on that new information.

The dissertation closes with a concluding section that will put focus on the potential fruitful avenues for further work departing from what is being developed here.

Chapter 2

Background: Justification, Fragmentation, Subjective Probability, and Social Choice

This chapter is devoted to providing an overview over the existing literature on topics relevant to this dissertation, as well as how this dissertation contributes to the literature on those topics.

I, first, give an overview over the concept of epistemic justification, its treatment through the history of philosophy, as well as an overview over the most prominent traditional contemporary accounts of epistemic justification.

Second, I introduce what I call the orthodox model of belief, that is, roughly, the idea that belief should be unified; that an agent's beliefs need to be consistent and closed under logic for the agent to qualify as rational. I comment on different possible interpretations of this orthodox model, and give an overview over streams in the more recent literature that diverge from the orthodox conception of belief in epistemology.

I, then, introduce fragmented epistemology as one particular way of diverging from epistemic orthodoxy, and more specifically, from the orthodox model of belief. I explain the central ideas of fragmented epistemology, provide an overview over motivations for fragmentation, delineate it from

contextualism as a very closely related framework (which is a discussion novel in the literature on fragmentation), and touch upon recent developments and remaining open questions in the fragmentation literature.

Since the following topics will be relevant for this dissertation, I also provide brief introductions into social choice theory and judgment aggregation, Bayesian epistemology, and imprecise probabilities. In the last section of this chapter, I go over what this dissertation will add to these respective streams of literature.

2.1 Justification

This section provides an overview over the literature on the topic of epistemic justification. I will first give a brief historical overview showcasing the central importance of the study of the concept of epistemic justification to philosophy generally, and modern epistemology, logic, and philosophy of science specifically. Then, I will delve deeper into the most prominent accounts of epistemic justification present in the orthodox epistemological literature, namely evidentialist accounts, reliabilist accounts, hybrid accounts, value property accounts, as well as contextualist accounts of justification.

2.1.1 Brief history

One can distinguish between different types of justification, and classify them broadly into epistemic justification and non-epistemic justification. This thesis will be concerned with epistemic justification – i.e. the sort of thing that is supposed to render a belief more rational / likely to be true. On the other hand, there are non-epistemic kinds of justification, ones involving non-epistemic or value-laden concepts such as morality ("Am I morally justified in doing/believing x / y / z?"). Those will not be the subject of this thesis.

The concept of epistemic justification has been of interest throughout the history of philosophy, and within contemporary analytic philosophy, is a central topic in epistemology, epistemic logic, and philosophy of science.

Within ancient and medieval philosophy, the concept of epistemic justification was discussed as so far as it was considered that one only had "scientia" or knowledge that something is the case if one understood *why* it is the case (cf. Swinburne, 2001). For example, in Meno, an early seminal work in epistemology, Plato writes that true opinion / belief, taken together with an explanation of that belief or opinion, is what constitutes understanding or knowledge (cf. Sedley and Long, 1980). What Plato does here, is to argue that true belief needs some kind of justification to constitute as knowledge. This exemplifies that, since the beginning, the interest in the concept of justification has been closely tied to (the study of) other important epistemological concepts, such as belief, understanding, and knowledge.

Within early modern philosophy, for example by Descartes, the concept of knowledge was understood in a different way – knowledge was tied to being infallible or absolutely certain (cf. Cottingham, Stoothoff, and Murdoch, 1984). And it was considered that, if one could not obtain knowledge that was absolutely certain and infallible, the only thing one could have is (justified) opinion or belief, but early modern philosophers were not very interested in that concept (cf. Swinburne, 2001), as knowledge was perceived to be the ultimate goal.

With the empiricist tradition, starting from the 17th and 18th centuries (such as for Hume), things changed. The idea that knowledge should be infallible was not supported anymore. The interest shifted from the concept of knowledge to the concept of belief, and empiricists thought that most beliefs are fallible / open to the possibility of error (Swinburne, 2001). If anything can be known, according to empiricists, it is the principles of logic and immediate sense data. Hume formulated the problem of induction on the basis of which it is easy to see how empiricists would reject the possibility of infallible knowledge. The problem of induction is about how one can justify an inference from singular statements to a universal statement (i.e. a hypothesis or theory). The problem is that there is no deductively valid inference from the fact that, for example, all (yet) observed ravens have been black, to the hypothesis: "All ravens are black". According to

the early modern conception of knowledge, an inductively generated statement, like the one about all ravens, would not classify as knowledge, since it cannot be deduced, and therefore cannot be known with certainty.

The problem that follows from this is that, if not this, then what can classify as knowledge? Not much, as in very many sciences and everyday circumstances, inferences with this inductive gap have to be made and are being made. Instead of the concept of knowledge, since knowledge can hardly ever be obtained, empiricists were more concerned with the notion of a justified belief, the concept of justification became important and required explicit analysis.

This leads me to discuss the concept of justification within the more contemporary analytic philosophy tradition.

2.1.2 Justification in contemp. analytic philosophy

It has been and still is one of the central goals of philosophy and epistemology of science to shed light on the concept of justification. One example of an account of justification within the philosophy of science literature is the hypothetical-deductive (H-D) model. According to that model, scientists (should[1]) develop a scientific hypothesis, and logically deduce testable empirical predictions from it. They then perform tests to see whether those predictions are true or false, which then either justifies or justifies to reject the hypothesis (this was endorsed eg. by Hempel, 1966 and Reichenbach, cf. Okasha and Thébault, 2020). This idea of justification is based on a generally deductive framework.

Very contrarily to this, Popper (1934) defended his famous falsificationism, where he argues that a scientific hypothesis can never be confirmed by data points, but only ever falsified. He thereby deemed the concept of the justification of a scientific hypothesis irrelevant – the only thing that counts is that whether the hypothesis has been defeated or not, and nothing else, specifically not whether a hypothesis or theory is justified, or true,

[1]This "small" word indicates a big philosophical discussion about whether models of justification should be descriptive or normative. I will delve deeper into this topic a bit later, especially when I talk about the orthodox model of belief.

or constitutes knowledge.

And then there is the Bayesianism as a theory of epistemic justification. Contrarily to the HD Model, and falsificationism, Bayesianism is an inductive theory of justification. This means, that it is not governed by deductive logical laws, but by the laws of probability. According to Bayesianism, the justification of a scientific hypothesis is explicated in terms of the degree of subjective probability an agent has in that hypothesis, where subjective probability is governed by two central tenets of Bayesianism – the Kolmogorov axioms of probability, and a conditionalisation tenet which tells us how to update our subjective probability in light of new evidence (Lin, 2024) (more on this in a bit). Within contemporary philosophy of science, Bayesianism is the most prominent account of justification.

Now, after having given a brief insight into the role justification plays within philosophy of science, let me show the role it plays in general contemporary epistemology and in logic[2].

I have presented how throughout history, different conceptions of the concept of knowledge have been defended. In the 20th century, a canonical definition of knowledge emerged within epistemology – the so-called "Justified True Belief" conception of knowledge. According to that definition, knowledge can be defined in terms of its constituents, a true belief that is also justified. The job of the concept of justification is to serve the role of a necessary condition for knowledge. This seminal conception of knowledge also showcases one of two main directions within contemporary epistemology – reasons-first epistemology. According to this direction, terms like belief and knowledge are complex and need to be defined in terms of other concepts that are conceptually prior and constitutive of those more complex concepts, such as reasons / justification (Schroeder, 2021). This dissertation is based on a reasons-first conception of epistemology, where the concept of justification is fundamental for definition more complex concepts, like belief and knowledge. Opposed to this is the Williamsonian knowledge-first epistemology, where knowledge, instead of justification is

[2]Though it is hard to sharply separate these discussions from each other. because they are closely intertwined.

taken as a primitive (T. Williamson, 2000) which will not be the focus of this dissertation[3].

As a response to the "Justified True Belief" definition of knowledge, Gettier (1963) published a short chapter where he develops counterexamples to that definition of knowledge. This article has sparked a whole industry of cooking up new definitions of knowledge and coming up with counterexamples to them.

But wait?! So far, the talk has mainly been of a definition of knowledge. What does this have to do with justification? Justification is being and has often been studied within its role as a necessary condition for definitions of knowledge. But justification is also being studied independently of the concept of knowledge, within philosophy of science, orthodox epistemology, and logic. And, I think, that concepts like justification (and belief) are especially worth studying, because they are (or can be) normatively demanding philosophical concepts without being so demanding, as knowledge often seems, that they are practically unattainable for real life agents. The concept of justification is being studied in epistemology, but also in logic, where logical systems are being created to formally represent the concept of justification (Artemov, Fitting, and Studer, 2024). This thesis will provide not a logic of justification, but an epistemological theory of justification that is formal in the sense that it uses formal tools to make the account more precise, but is not a logic.

Before developing a new account of justification, it is important to provide an overview over the common accounts of justification already present in the epistemological literature[4].

[3]For some arguments for why reasons-first epistemology might be superior to knowledge-first epistemology, refer to Gerken (2018) or Schmidt (2020).

[4]Since my thesis is within the field of epistemology, I focus on accounts of justification from orthodox epistemology strictly, and do not elaborate further on accounts of justification from philosophy of science or logic (with the exception of Bayesianism).

2.1.3 Traditional accounts of justification

I now present an overview over existing accounts of justification in the epistemological literature.

All of the accounts I will present are predominantly synchronic accounts of justification. This means that they are predominantly concerned with "what it is for a belief to constitute a justified response to the [epistemic] situation in which the believer finds herself at a given time" (Swinburne, 2001, p. 1f.). They are not diachronic accounts that would focus on what it means for a belief to be justified through time, how justification might be revised in light of new information, and in dynamic settings[5].

Before presenting the different accounts, it will be useful to present some concepts that will help to classify and distinguish them better.

First, the literature on justification distinguishes between internalism and externalism about justification (cf. eg. Swinburne, 2001, Fumerton, 2002, Steup and Neta, 2020)[6]. Broadly speaking, internalism about justification captures the idea that what is relevant to determine justification is entirely internal (i.e. amounts to mental states or experiences). This means that one analyses the concept of justification in terms of factors internally accessible to an agent, for example in terms of other (types of) beliefs of the agent.

On the other hand, externalist theories of justification analyse justification in terms of factors which might be mentally inaccessible by an agent, such as whether a belief is produced via a reliable external process (more on this in a bit). The account of justification that will be developed in this dissertation will be internalist.

Moreover, another relevant distinction within the taxonomy of theories of justification is that between foundationalism and coherentism about justification. Foundationalism about justification is the view that justified

[5]The latter chapters on fragmented Bayesianism in this dissertation do talk about some dynamic aspects of justification. But this is not the focus of this dissertation, and would be a potentially fruitful future research avenue.

[6]This distinction does not only play a role concerning accounts of justification, but is also relevant in other areas of philosophy, such as philosophy of mind.

beliefs / the justification of belief is structured like a building; they are divided into a foundation (consisting of basic beliefs which are not justified by any other beliefs) and a superstructure (consisting of non-basic / inferentially justified beliefs). The main objections to foundationalism are, first, the denial of the existence of basic statements, and second, the claim that, even if basic statements exist, they are too weak of a basis to justify all the other statements in the superstructure (cf. Annis, 1978). Of course, one might ask many further questions here, which are discussed within the literature about foundationalism, such as: "What should / does an account of basic belief look like?", or "What is / should be the logical relationship between the base and the superstructure?" Answering those questions is not part of the focus of this dissertation.

Coherentism about justification, as opposed to foundationalism, is the view according to which justified belief is structured like a web, and the justification of belief should be understood in terms of the way in which the proposition of interest coheres with other propositions (either by the same agent or by a community of agents). Again, there a many open issues. For example, it is prima facie not entirely clear how the concept of coherence should be spelled out precisely. Consistency, for example, is sometimes argued to be too weak of a criterion, since the web of beliefs entertained by an agent could have nothing to do with each other in terms of the content of the propositions, but nevertheless be consistent. This would seem insufficient, because why should a proposition about arithmetic be justified via it being consistent (i.e. non-contradictory) with a proposition about football? This example shows that, intuitively, mere consistency with other beliefs seems insufficient to justify any given belief[7].

[7]Foley (1979) even argues that consistency might, in fact, not be a necessary condition for justified belief because of Lottery-type situations (for an overview, cf. i.e. Hawthorne, 2003): Imagine a Lottery with 1000 tickets. Each person buying a ticket might be justified in believing that they lose. However, at the same time, each person might also be justified in believing that there is a winning ticket. But, by closure under conjunction, each person's justified belief that they will not win taken together is inconsistent with the justified belief that there is a winning ticket. The two obvious solutions to this paradox are to either give up closure under conjunction or to admit that justified beliefs do not have to be consistent with each other, and thereby deny consistency as

The account of justification developed in the first qualitative part of this dissertation will be compatible with both, a version of foundationalism as well as a version of coherentism about justification. It will be compatible with foundationalism, because one will be able to perceive the account developed here as allowing that there will be more "fundamental" propositions called pieces of evidence that can play the role of standing in a justifying relation to another, less "fundamental" proposition in need of justification.

The account will also compatible with a kind of local coherentism. It will not be the agent's global belief state that needs to be coherent, but instead, fragmented justification will require a kind of "local" coherence. It will become clear throughout the dissertation what this means exactly.

Evidentialism

Evidentialism is often thought of as the default or common sense account of justification. The idea is that justification is to be spelled out in terms of having good reason for a belief.

There being many different options for precisely spelling out this account, a standard formulation of evidentialism can be found in Conee (2004):

> "Doxastic attitude D toward proposition P is epistemically justified for S at time t if and only if having D toward P fits the evidence S has at t."

Traditionally, evidentialism has been conceived as an internalist account of justification, but this view has been increasingly challenged, with arguments that evidentialism is also compatible with externalism, following a development in the conception of evidence (from private to public and objective) (cf. Kelly, 2016).

Furthermore, evidentialism has been argued to fit more naturally into a coherentist picture about the structure of justification, since coherentism

a necessary condition for justified belief. Note that Leitgeb (2017) provides a solution to this paradox, where, due to a contextualist argument, neither consistency nor closure under conjunction must be given up.

allows for evidential support for all of one's justified beliefs, whereas foundationalism requires some beliefs to be justified simpliciter (i.e. justified without requiring any further evidence as justification). However, evidentialism might also be compatible with foundationalism, if one allows the pieces of evidence in the base structure that support a belief in the superstructure to qualify as beliefs themselves, and interprets the difference between beliefs in the base and the superstructure as being supported by different types of evidence.

The account of justification developed in the first qualitative part of this dissertation will be broadly evidentialist, though not purely. Evidence will play an important role in whether a proposition is justified. But, what will be even more important is the *relation* in which that evidence stands towards the proposition in question. Within an evidentialist account of justification, it is important to properly spell out the concept of evidence, since it is the main concept involved in defining justification. Within my account, evidence will be modelled via propositions (i.e. sets of possible worlds; more on this shortly).

One problem with evidentialism is that having evidence in support of a believed proposition is a necessary but insufficient condition for that belief to qualify as justified. This is because not all evidence is equal. Evidence might be untrustworthy, i.e. because it comes from an unreliable source, and therefore be unable to render a belief justified. For example, evidence coming from an unreliable / biased / (government-)controlled news channel might not be the best basis for a justified belief. This is a deficit that the next account of justification aims to tackle – reliabilism.

Reliabilism

Another prominent account of justification is reliabilism - of which one of the most cited versions was developed by Goldman (see eg. Goldman, 1979) and is called process reliabilism. According to process reliabilism, whether a belief is justified depends on the "reliability of the process(es) that cause it, where reliability consists in the tendency of a process to produce beliefs

that are true rather than false" (Goldman, 2012, Ch.2). But how reliable must a process be in order to justify the beliefs it brings about? Perfect reliability seems too high of a standard, since, intuitively, belief formation processes which do err but do so very rarely can still confer justification. On the other hand, the required degree of reliability should not be too low, either. If it were less than 0.5, them it would err more often than not, and such a process could hardly confer justification. Therefore, the required degree of reliability should lie somewhere between 0.5 and 1. Any further general specification of the exact level seems ad-hoc, and it has been argued that the precise level of justification required for a process to confer justification is context-dependent (Steup and Neta, 2020).

Reliabilism is perceived as an externalist theory, since what is important, here, is whether a belief is likely to be true given a certain belief-formation process, and moreover, since processes are not necessarily internal mental states.

Process reliabilism faces a number of issues, such as, among others, the Clairvoyance Problem, the Generality Problem or the Problem of Defeat. A majority of the issues reliabilism faces, including the three ones just explicitly referred to, are due to a lack of the core evidentialist intuition within the reliabilist theory of justification. To keep things concise, I will exemplify this using the Clairvoyance Problem. This problem, first formulated by Bonjour (1980) asks one to imagine a clairvoyant called Norman who has a perfectly reliable clairvoyance faculty. But, he has no evidence or reason for or against the general possibility of such a power, or for or against his possessing such a power. One day, his clairvoyance faculty generates the belief that the US president is currently in NYC. But there is no evidence for or against this belief. BonJour argues that, intuitively, his belief is not justified, since it lacks being based on some kind of evidence or reason, but according to reliabilism, the belief would be perfectly justified.

Hybrid accounts

Evidentialism and reliabilism have often been viewed as incompatible rivals, one theory being an internalist and the other an externalist one. But, evidentialism has increasingly been argued to be compatible with externalism (see eg. T. Williamson, 2000, Bergmann, 2018). Furthermore, each position seems to fall short of capturing an intuition about the concept of justification that its rival is able to capture. Recently, quite a few philosophers, such as Comesaña (2010), Goldman (2011), Goldberg (2018), E. Miller (2019), Pettigrew (2021), have proposed hybrid accounts of the two and argued for their compatibility. All of these hybrid accounts share the same idea: The definition of justified belief has two necessary and jointly sufficient conditions, a reliabilist and an evidentialist one.

To give an example of one such hybrid account, I choose Goldman (2011)'s formulation of a hybrid view, which is called the "Two-Component Hybrid View": Agent S's belief that P is justified if and only if both:

1. S's total evidence at time t supports believing P, and

2. S's belief that P is the result of a reliable belief-forming process

The account developed in the first qualitative part of this dissertation could be viewed as a hybrid account, as it will rely on evidence as a necessary condition for justification, but also will require that evidence to stand in a valid justification relation towards the proposition in need of justification in question. A relation is not a process, but similar to hybrid views, my account will require a valid relation between evidence and proposition to be justified, in addition to mere evidence.

Value property accounts

Another prominent set of accounts of justification can be called "value-property" accounts. According to such accounts, a proposition is more or less justified, and it is not specified what entity a justification amounts to. This group of accounts can also be called quantitative or probabilistic, as

opposed to the accounts presented above, which are all qualitative accounts of justification. Here, what exactly justification "is", is left covert, and a proposition gets assigned a justification value, similar to a truth value. The values can be zero and one, for "justified" and "not justified", but they can also cover the real interval between zero and one, as it is the case in Bayesian epistemology (Lin, 2024).

For example, in her paper on justified group belief, Lackey (2016) uses such an account. A very prominent value property account of justification is the Lockean account of justified belief (Foley, 2009, Leitgeb, 2017). According to the Lockean account of justified belief, a belief is justified (simpliciter) if the doxastic agent has a sufficiently high degree of belief in that proposition. A threshold determines whether the degree of belief in a certain proposition is high enough to qualify as a justified belief (simpliciter), or not. Often, the factor determining the value of this threshold is coined "context", rendering the justification of belief a context-sensitive notion. This Lockean account account is a value property account, because it masks what kind of information the doxastic agent needs to possess and collect in order to reach a degree of belief higher than the Lockean threshold, and thus in order to render her belief in a certain proposition justified. It is not explicitly spelled out in these accounts whether what boosts the agent's degree of belief, and eventually lifts it over the threshold is (propositional) evidence, a reliable belief formation process or something else. Instead, what does the work here is probabilistic Bayesian machinery. I will explain more about Bayesianism in section 2.5 of this chapter, as well as the 5th and 6th chapter of this dissertation, since they will deal with a fragmented account of Bayesianism, and thus deal with a fragmented version of value property account of justification.

Contextualist accounts

Then, there are contextualist accounts of justification. Conceptually, they are on a different level from the accounts presented so far, because each, evidentialism, reliabilism, hybrid accounts, or value property accounts can

be – hypothetically[8] – combined with a contextualist understanding of justification.

Furthermore, it is important to present contextualist accounts of justification, because contextualism and fragmentation are closely related, and it is therefore necessary to make explicit how existing accounts of contextualist justification differ from the new account of fragmented justification I am developing in this dissertation.

First, let me draw a distinction between two types of contextualism, (proper) semantic contextualism, as opposed to a non-semantic contextualism. This distinction is important, because existing accounts of contextualist accounts of justification are of the first kind, whereas fragmentation is more closely related to (but not identical to) the second kind. The first, proper, notion of contextualism is semantic; according to it, the content of the concept in question, such as justification, changes from context to context. Here, the mental state of the agent that is reasoning in question does not have to change – what changes, here, is the content of the concept, but not necessarily the mental state of the reasoner. This kind of contextualism is often described in terms of taking the perspective of belief ascription / attribution. The choice of content / truth conditions of a concept like justification, here, depends not on the mental state of the reasoner, but on features of the situation in which the belief is ascribed (cf. Leitgeb, 2017).

The second non-semantic notion is different; here the content of the concept of justification does not change from context to context, but remains fixed. Rather, what changes is the agent's non-epistemic and epistemic circumstances (such as their stakes, perspective, doxastic state,...), and therefore the truth value of a proposition containing the concept in question might change in different contexts, without the truth conditions / content of the concept changing. Here, the perspective is not that of someone who

[8]Annis (1978) is the first to develop a contextualist account of justifiation. Annis presents his account as an alternative to existing foundationalist and coherentist accounts, but the account is not per se incompatible with either evidentalism, reliabilist, hybrid accounts, or value property accounts, but adds a contextualist dimension to them. Annis intends contextualism to be another variant of possible type of justification – an account of justification can be foundationalist, coherentist, or contextualist.

ascribes or attributes for example a belief or justification to someone else. Rather, the perspective in question is the mental state of the agent under consideration – whether or not an agent will have justification will depend not on an external ascriber, but on the epistemic and non-epistemic aspects of the mental state of the agent themselves. It, thereby, is an internalist notion of contextualism.

This second type of contextualism will be important, because it is very closely related to fragmented epistemology. This type of contextualism (and variants of it) has received different names in the literature, Leitgeb (2017) calls it non-semantic contextualism, Hawthorne (2003) calls it sensitive moderate invariantism, Stanley (2005) calls it interest-relative invariantism. It is also often found in connection with the literature on pragmatic encroachment (Fantl and McGrath, 2009)[9]. Pragmatic encroachment is the idea that unlike epistemic orthodoxy prescribes, the pragmatic realm does (and maybe should) encroach on the epistemic realm. According to orthodox epistemology, what counts as knowledge, belief or justification solely depends on epistemic factors. Non-epistemic, i.e. practical, moral or other contextual factors ought not influence what counts as knowledge, belief or justification. And this is exactly what non-semantic contextualism and pragmatic encroachment argue against. Practical factors within a context, such as stakes, perspective, interest, but also epistemic factors within a context, such as the doxastic state of the agent, their background knowledge etc., can influence what counts as knowledge, belief or justification in a certain context, according to those accounts.

After giving an overview over these two distinct types of contextualism, let me now go on to present the seminal account of contextualist justification (more precisely, justified belief). I will show how this account is of the first type of contextualism just explained – semantic contextualism, instead of the non-semantic kind that will be closely tied to fragmentation.

A contextualist account of justified belief has first been developed by Annis (1978), with further developments and variations by S. Cohen (1999),

[9]Cf. also Thomason (1986), Nozick (1994), Bratman (1999), Thomason (2001) and Cresto (2010).

and Wedgwood (2008) (cf. Rysiew, 2021). Note that strictly speaking, those are contextualist accounts of justified belief and or knowledge, and not the concept of justification in isolation. Unlike these, my account of fragmented justification will provide an account of justification within the broadly contextualist spirit that it strictly about justification, and not belief and / or knowledge.

Annis presents his contextualist theory of justification as a superior alternative to the existing types of accounts of justification; foundationalism and coherentism. According to Annis, what is missing from both foundationalist and coherentist accounts of justification is the fact that justification has an important social component. According to Annis' contextualist account, an agent is justified in believing a proposition P only if, relative to the issue context which determines the level of understanding and knowledge required, and an appropriate objector group, the agent can satisfactorily meet objections of that objector group. The objections have to express a "real doubt", and the objector group has to consist of "critical truth seekers". "Social information – the beliefs, information and theories of others play an important part in justification, for it in part determines what objection will be raised, how a person will respond to them, and what responses the objectors will accept" (cf. Annis, 1978, p. 215).

This account is contextualist in the sense that, in different issue contexts and depending on the topic under consideration, the relevant objector group will be different, the relevant objections might be different, and the way the agent has to respond to them will be different. Imagine the proposition in question being about a the causes of a rare type of cancer. Depending on the context of issue, the kind of justification the agent has to put forward to defend that proposition will be very different, according to Annis' theory. Imagine one issue context to be a primary school, and another issue context to be a board of medical experts on rare types of cancer. Depending on the issue context, the appropriate objector groups will be different, according to Annis. In the first issue context, the appropriate objector group will be other primary school teachers, maybe even primary school children. In the second issue context, the appropriate objector group

will be medical experts and researchers on rare types of cancers. One can easily imagine how the types of questions asked and objections raised in the different context will be very different. Depending on the issue context, the topic, the relevant objector group and their demands, what will count as an appropriate justification will differ vastly.

Let me explain how this theory of contextualism is of the first of the two kinds presented earlier. What determines whether an agent is justified or not in believing a proposition is not (primarily) the agent themselves, but instead, it is the objector group and whether they deem the agent to appropriately respond to the objections and challenges they brought forward. Hence, this is a kind of attributer contextualism (that is, semantic contextualism) and not reasoner / non-semantic contextualism. The objector group as an objective, context-dependent social standard decides whether or not the agent in question is actually justified in believing a proposition or not. This type of contextualism is semantic; the content of the concept of justification changes depending on the context. The type, number and quality of reasons an agent has to put forward to appropriately justify a proposition vastly depends on the context, the topic and the objector group. It is not just the truth value of a justification statement that changes from context to context (which would amount to non-semantic contextualism), but rather what constitutes the concept of justification can fundamentally change from context to context.

While this type of contextualism is interesting, what is lacking from the literature is a contextualist account of justification in the second non-semantic sense, where the content of the concept remains invariant across contexts, but rather, where, what changes is the epistemic and non-epistemic circumstances of the reasoner themselves. Fragmented justification will be a version of such an account. Of course, one might immediately ask, what the difference between a non-semantic contextualist account of a doxastic attitude versus a fragmented account of that attitude is. This is an important question I will answer in section 2.3.5 of this chapter, after having introduced fragmented epistemology.

Fragmented epistemology has come about as a response to the concep-

tion or model of belief within orthodox epistemology which I will talk about now.

2.2 The orthodox model of belief

The orthodox model of agents' belief states is often described using the metaphor of a web (cf. Bendaña and Mandelbaum, 2021) or map (Ramsey, 1931). Characteristic for this orthodox conception of agents' belief states and their rationality is that they should be unified in some way. Bendaña and Mandelbaum (2021) describe this need for unity of the state by using a web metaphor. For an agent to be rational, their doxastic state should be interconnected in one big web-like network, where the entire web guides reasoning and action. In similar spirit, much earlier, Ramsey (1931) describes belief as the "map by which we steer". This map needs to satisfy certain rationality criteria to be able to sensibly guide reasoning and action. Kindermann and Onofri (2021) summarise that "much of epistemic logic, Bayesian accounts of belief, decision theory and some influential views about rationality proceed from the underlying view that the mind – at least doxastic states – is unified". They describe the requirement of unity as a conjunction of two theses about agents' doxastic states. (Rational) agents need to have a unified mental representation of the world (at a specific time t[10]), that is, a single belief state (at t) that is organised by two principles (cf. eg. Quine, 1951, Cherniak, 1983, Fodor, 1983):

1) Consistency: The total set of an agent's beliefs (at t) is consistent.

2) Closure: The total set of an agent's beliefs (at t) is logically closed. That is, an agent believes all logical consequences of their beliefs.

This means, for example, that even beliefs that seem "deeply unrelated would count as connected to each other through consistency" (Bendaña and Mandelbaum, 2021), such as, for example, beliefs about cats and beliefs

[10]Recall that I am focusing on synchronic justification in this dissertation, as opposed to diachronic justification.

about mathematics. Furthermore, if all beliefs are connected with each other like this, new information that affects the status of one single belief within that unified structure can have consequences for the entire structure (this feature is sometimes referred to as the Quinean nature of belief; cf. Fodor, 1983).

Slightly more detailed than Kindermann and Onofri (2021)'s summary, Bendaña and Mandelbaum (2021) summarise the orthodox model of belief as the conjunction of the following theses:

1) Unity: Beliefs are stored in one single database.

2) Quineanism: Belief revision is sensitive to global properties of an agent's total set of beliefs.

3) Consistency: When any of an agent's beliefs change, all other beliefs adjust to remain consistent with the modification.

4) Uniqueness: Belief storage does not contain redundant token representations.

5) Conservatism: The more revisions a belief change would require, the less likely it will be that that change will occur.

Unity is defined in a different way here than above. Here, it simply states that beliefs are structured as one big web or map, not multiple different ones. The consistency requirement is pulled apart from the unity requirement in this way of formulating the orthodox model of belief. Quineanism amounts to the thesis already mentioned that whenever an agent acquires new information, the belief revision process that happens to accommodate this new piece of information happens on the global doxastic state of the agent. Consistency is clear and similar to the consistency requirement stated above. Uniqueness and Conservatism are new. I will briefly explain them, but they are not part of the primary focus for this dissertation. Uniqueness is the thesis that each belief may only be represented once in a doxastic state. This means that the proposition P expressing the belief

that the earth is round may only be stored / represented once in an agent's belief state, not multiple times, i.e. redundant token representations of a belief type are not allowed. The last item on this list – conservativism – has a different flavour than the previous ones. Theses 1) to 4) are have a normative flavour, but conservativism seems to aim to descriptively capture a fact about the psychological reality of belief states of agents. Bendaña and Mandelbaum (2021) argue that because of conservativism and consistency, "people will be less likely to revise beliefs about logical or mathematical truths than beliefs about nearly anything else. This is because changing logical and mathematical beliefs entails a potentially infinite number of changes to the web" (p. 80). This seems more of a descriptive or pragmatic / ecological thesis than a normative one. Within the psychology of agents' beliefs, some beliefs are more tightly held onto than others, and those that are held most tightly onto are those that would require the most work to get rid of via the belief revision process.

This hints at a very important discussion to be had about the orthodox model of belief – is it intended to be a normative guide for the rationality of agents or is it intended to be a descriptive account of (real or ideal) agents?

There are at least three ways the orthodox model of belief could be interpreted, as a descriptive model of real agents, as a descriptive model of ideal agents, or as a normative guide for the rationality of real agents.

The orthodox model of belief would be a very bad descriptive model of the psychological reality of real agents' doxastic states. There is plenty of empirical evidence that provide counterarguments against the web model being an accurate representation of how real agents reason (cf. Porot and Mandelbaum, 2020, Bendaña and Mandelbaum, 2021). Actually, this evidence that I will provide an overview over shortly, strongly supports a fragmentation model of belief. Real agents do not and cannot live up to the strong requirements that the orthodox model if belief imposes. Their beliefs contain inconsistencies, redundancies, and they are not closed under logic, because real agents are not logically omniscient, but highly fallible and limited. In the literature, this is called the problem of logical omni-

science[11]. Understood as a descriptive picture of real, limited, and fallible agents, this orthodox model of belief is highly inadequate. It would greatly miss its goal of representing real agents' doxastic states if understood like this.

However, one might argue that it is in the nature of the orthodox model being a model that it is impossible for it to be fully descriptively adequate representation of the psychological reality of agents. One might argue that it, in fact, does not have to be fully descriptively adequate: It is in the very nature of a (scientific) model that certain aspects of reality have to be idealised away, otherwise it would not be a model. There is plenty of research in the philosophy of science literature on the topic of models (in science), what they are, what they represent and how they represent, how they explain and how they use idealisations to achieve modelling goals (cf. eg. Bokulich, 2011, Frigg and Nguyen, 2019, Frigg and Nguyen, 2021, Frigg and Hartmann, 2024).

Models, in order to serve their purpose of explaining, predicting, fostering understanding and so on, have to idealise and distort some aspects of their target system, for various reasons, such as simplicity, computational tractability, but also because some aspects of the target system might be irrelevant for studying a particular phenomenon within that system (Aristotelian idealisations), or because the distortion might not be faithful to reality of the target system, but might still be useful in the sense that one can understand the idealised model as the ideal limit of the real system (Galileian idealisations). One might argue that the way to understand the orthodox model of belief is by appreciating that it involves significant idealisations with regards to the reasoning abilities of real agents.

In fact, the orthodox model is often conceived to be a descriptive theory of idealised agents with perfect rationality (cf. Stalnaker, 1991). The problem with this interpretation of the orthodox model is how useful such a descriptive theory of ideally rational agents is. Often in the literature, this theory is intended to describe idealised versions of real agents. This

[11]Cf. eg. Stalnaker (1991) for a presentation of the problem of logical omniscience.

is supposed to still be informative about real agents by providing a model describing the ideal limit of real agent's doxastic states. Kindermann and Onofri (2021) summarise that it is an open question "whether real agents are similar enough to these idealised agents for the theory to have any use in the explanation of real agents' doxastic attitudes".

Of course, still, a proponent of this reading of the orthodox model of belief might argue that it is useful and interesting in its own right to study rationality criteria for ideal agents that real agents might never be able to fulfil. However, because of the motivation of the project of dissertation I laid out previously, I am interested in the application of epistemological theories to real agents, by understanding those theories as a normative guide for the rationality of real agents, which is the third possible reading of the orthodox model of belief.

The third option to understand the orthodox model is as a normative guide for the rationality of real agents. This avoids the problem of lacking descriptive adequacy of the model, because, understood normatively, descriptive adequacy is no longer one of the goals of this model. But this third understanding of the model is problematic for its own reasons. If understood this way, it provides standards of rationality that are unattainable by real agents. Real agents, even the most epistemically able ones, such as scientific experts, will fall short of the demands posed by the orthodox model. No belief system by a real agent is closed under logic; real agents, even if highly epistemically able, are limited and fallible.

There are two possible viewpoints one can take in regard to this. One can argue that this is not a problem, since the model is a normative guide that does not necessarily have to be achieved by real agents. But the second possible viewpoint is that, then, no real agent would ever count as rational relative to the standard imposed by the orthodox model. This could be problematic, because, intuitively, there are many real agents whose belief systems seem rational although they do not adhere to the strict standards imposed by the orthodox model. Because no real agent can adhere to the standards set by this model, the standards of rationality postulated by this model might be too strong to be useful in deciding, in practice, which agents

are rational and which ones are not. No real agent would classify as rational according to the normative reading of this model, which seems very strong, given that, intuitively (and this intuition is not captured by the orthodox model), it seems that many agent exhibit some kind of rationality although they fail to adhere to the norms postulated by the orthodox model of belief.

Again, one can argue against this position, by saying that it is not the goal of a normative account of rationality to allow for real agents to actually count as rational. But the kind of epistemology that I am most fascinated by is epistemology which aims to be a normative guide for real agents but without the concept in question – such as rational belief – being completely unattainable in real life.

Not surprisingly, it is not only me who is interested in this kind of epistemology. There are many different approaches of weakening the orthodox model of belief understood as a normative guide for the rationality of real agents, to be demanding enough to be normatively significant, but "weak" enough to allow for real agents to actually attain this standard. I will now give an overview over such approaches, before zooming in on fragmentation as one particular such approach.

Some authors challenge the consistency requirement on rational belief, and propose weaker notions to use instead of it. For example, Fitelson and Easwaran (2015) introduce a notion called non-dominance which is strictly weaker than consistency to replace the consistency norm . But, from the two requirements posited in "Unity" above, consistency is the more challenging one to let go of, since it is often conceived to be a minimal requirement on rationality or rational belief (cf. Davidson, 1973, Davidson, 1982a, Davidson, 1982b, Kolodny, 2008, Fitelson and Easwaran, 2015). Instead, more authors choose to give up or weaken closure under logic, a requirement which brings about the problem of logical omniscience (Stalnaker, 1991). Alternatively, one can give up the framework of classical logic for rational belief altogether and move to non-classical logic approaches, eg. by dropping monotonicity, closure, or the consistency requirement (cf. Artemov, Fitting, and Studer, 2024, Strasser and Antonelli, 2024).

Recent examples of responses to the problem of logical omniscience are

(but by no means limited to) Bjerring, 2013, Smithies (2015), Skovgaard-Olsen (2017), Dogramaci (2018), Skipper and Bjerring (2019), and Hawke, Özgün, and Berto (2020). Savage (1967) introduces the problem of logical omniscience for credence, responses / solutions include De Finetti (1970), Good (1983), Schervish, Seidenfeld, and Kadane (2002), and Konek (2022).

Recent work in epistemology surrounding formulating (normative) theories for less than perfectly rational agents includes (but is by no means limited to) Gigerenzer (2006), Fantl and McGrath (2009), Icard (2018), Staffel (2018), Dorst (2019), Douven (2020), Dorst and Mandelkern (2021), Barrett (2022), Dorst (2022), Greco (2023), and Thorstad (2024). I am deliberately omitting citing the research paths on contextualism and fragmentation here, as I will turn to these in the next section.

Similar trends can also be recorded within the social epistemology (for an introduction cf. O'Connor, Goldberg, and Goldman, 2024, Goldman, 2005) and agent based modelling literature. One of the central goals of social epistemology is, contrary to orthodox epistemology, to study "how people can best pursue the truth with the help, or sometimes in the face of, other people or relevant social practices and institutions" (O'Connor, Goldberg, and Goldman, 2024). It is concerned with how individual agents can acquire things like justified belief and knowledge in a social, cultural and institutional context. Furthermore, it is also concerned with the application of central questions in orthodox epistemology to groups and collective agents. Central for pursuits like this is to diverge from the strong normative ideal that orthodox epistemology imposes, and accept that social and cultural factors, and things like biases and the like, can and do influence the pursuit of justified belief. But, the aim is not to be fully descriptive here either, but to accept that certain (potentially non-ideal) factors influence epistemological pursuits for individuals and groups, and come up with models of this that are less normatively demanding than orthodox epistemology, but not fully descriptive either. Notably, social epistemology finds that things one might think hinder the pursuit of belief or knowledge,

might in the end aid it (eg. Gabriel and O'Connor, 2024[12]).

In practical rationality, such as decision theoretic Bayesianism, expected utility theory, and economics, there are similar trends. Simon (1955) / Simon (1957) first introduces the term "bounded rationality" within the context of economics and expected utility theory. His idea is to "replace the perfect rationality assumption of homo economics with a concept of rationality better suited to cognitively limited agents" (Wheeler, 2024). Traditionally, expected utility theory assumes perfect rationality, this means "complete information about the options available for choice, perfect foresight of the consequences from choosing those options, and the [ability] to solve an optimisation problem that identifies an option which maximises the agent's personal utility" (Wheeler, 2024).

Expected utility theory provides an axiomatic formulation of what a rational agent's preference relation over her options / choices needs to look like (cf. Bernoulli, 1738/1954, Von Neumann and Morgenstern, 1944). Here, bounded rationality literature takes some of the assumptions on the preference relation, and weakens them to account for less than perfectly rational economic agents[13].

This brief overview shows that there is much interest in and research about responses to orthodox epistemology. Not among the overview just provided was fragmentation. Fragmentation is another, fairly recent, response to orthodox epistemology, the orthodox model of belief and its highly normatively demanding standards for agents which is at the centre of this dissertation.

[12]Often, work in social epistemology is done via so-called agent based models, computer models of epistemic agents (in social settings); for an overview, see Šešelja, 2023

[13]Examples of this are denying the preference ordering to be complete and / or transitive, eg. Keynes (1921), Koopman (1940), Luce and Raiffa (1958), or Aumann (1962). The other axioms can also be weakened / dropped, eg. by denying the independence axiom on the preference ordering, because empirical research suggests that agents evaluate gains and losses differently, cf. Friedman and Savage (1948), Kahneman and Tversky (1979).

2.3 Fragmented epistemology

As already mentioned, fragmented epistemology and fragmented models of belief are another reply to orthodox epistemology and orthodox models of belief that demand something like the principle of unity presented earlier.

Fragmented epistemology, of which the earliest accounts were developed by Davidson (1982b), Lewis (1982), Cherniak (1983), Stalnaker (1984), and Cherniak (1986) do not demand unity of an agent's epistemic or doxastic state as constitutive of rationality generally, and rational belief specifically. Fragmented epistemology allows for the belief state of a (rational) agent to be fragmented – i.e. compartmentalised into different sub-states or subsets of the total belief state. And crucially, fragmentation accounts relax the strong rationality norms imposed by orthodox epistemology. Within each fragment in the agent's doxastic state, the orthodox rationality norms of consistency and closure under logic remain present, but across fragments, they do not have to hold. Across fragments, inconsistencies are allowed and closure under logic does not need to hold.

The general idea of fragmentation views is to keep the central idea of orthodox epistemology that consistency and closure are important components of rational belief (cf. Kindermann and Onofri, 2021) while avoiding many problems that orthodox epistemology faces, by reducing the rationality standards for agents by allowing for fragments that might be incompatible with each other in certain respects.

A very important aspect of fragmentation accounts, again, is the question whether this is a descriptive or a normative account.

Many authors make it seem like fragmentation accounts are descriptive accounts of how real, fallible agents reason in real life. For example, Egan and Davies (2013) describe fragmentation like this: "Actual belief systems are fragmented and compartmentalised. Individual fragments are consistent and coherent, but fragments are not consistent and coherent with each other and different fragments guide action in different contexts. We can hold inconsistent beliefs and act in some contexts on the basis of the belief that P and in other contexts on the basis of the belief that $\neg P$. Frequently,

we fail to put things together or to "join up the dots". It can happen that some actions are guided by a belief that "If P then Q", but no action is guided by a belief that Q because the belief that P and the belief that if P then Q are in separate fragments" (Egan and Davies, 2013, p. 705) and inferences can only be drawn within, but not across fragments.

This makes it seem like fragmentation is a descriptive account. Real agents cannot live to the norms of orthodox epistemology, but fragmentation is not supposed to replace orthodox models of belief as a normative guide for the rationality of agents, but rather, it is intended to be a (supplementary) descriptive account of how real agents actually reason.

I will provide arguments to the contrary – I will argue that fragmentation accounts should be accompanied by rationality norms, and they then can and should be viewed as normative accounts of the rationality of agents. Fragmentation is not just something agents only do because they cannot do better. I argue that, in some epistemic circumstances, for specific reasons which I will delve into in this dissertation, fragmentation is the epistemically superior strategy compared to the strategies that orthodox epistemology provides. Fragmentation can be the rational thing to do sometimes. So, I will interpret and present my account of fragmented justification not as a descriptive account, but a normative one, albeit normatively weaker than orthodox epistemology demands.

Of course, it requires argument for why an epistemology that is so much "weaker" than orthodox epistemology can be a good guide for rationality. This is one of the most debated topics in the literature surrounding fragmented belief (cf. eg. Borgoni, 2021, Yalcin, 2021), and I will add to this literature by providing a list of reasons for why fragmentation (and specifically fragmented justification) can be rational, by developing rationality norms for fragmentation, and by showing how fragmented justification boosts the attractiveness of fragmented epistemology generally, because it provides a rationality argument for fragmented belief. Fragmentation, in my view, is not an account of describing or rationalising epistemic mistakes, but fragmentation can, in some cases, be the epistemically superior thing to do.

There is another important thing to note from the quote by Egan and

Davies (2013) just touched upon. They note that different mental fragments of an agent are somehow connected to different practical contexts and tasks. They say that different fragments guide action in different contexts. And this is why, according to many, the failure of fragmented agents to globally, inter-fragmentarily, adhere to consistency and closure is not that bad, because it is only ever at most one fragment that guides action in a specific context but not more. It is never the entire (potentially inconsistent) belief state of an agent that guides action, and therefore, facing practical decision scenarios, the agent never has to face their inconsistent beliefs in their entirety and, thus, this does not pose a problem for practical rationality (cf. also, eg. Elga and Rayo, 2021, Kindermann and Onofri, 2021, Elga and Rayo, 2022).

This is another point where I will depart from the existent literature on fragmentation. I aim to pull apart epistemic from practical rationality in the context of fragmentation. I take fragments to not necessarily be directly tied to a particular practical task (thus, I provide a new way of individuating fragments), and thereby to practical rationality, and I allow for multiple fragments to become relevant for one and the same practical decision. The account of fragmented justification I develop is primarily epistemic rather than practical in nature. I will argue that, purely epistemically, it can be the case that a fragmented mental state is rational. Within my framework, it can be the case that epistemic and practical rationality work in unity with each other (when a fragmented belief state poses no problems for practical action), but they can also stand in tension with each other. It can happen that multiple different fragments will become relevant to a certain practical task or decision. If those fragments that need to be active for a certain decision contain inconsistent information, then fragmentation (although epistemically rational) will pose a threat for practical rationality. Situations like this will make it necessary to think about how a process of de-fragmentation can look like in case the fragmented state poses threats to practical rationality which is a relatively undiscussed topic in the fragmentation literature. I will introduce mechanisms for de-fragmentation from the judgment aggregation literature, and will eventually argue that

de-fragmentation is difficult, and that, often there will be no epistemically optimal aggregation mechanisms. I will provide an argument along the lines of the seminal inductive risk argument by Rudner (1953) for the necessity of the employment of non-epistemic value judgments in the de-fragmentation process.

I will now present the origins of fragmentation, by summarising both the Cherniak (Cherniak, 1983, Cherniak, 1986) as well as the Stalnaker / Lewis model of fragmented belief (Lewis, 1982, Stalnaker, 1984). Then, I synthesise the central ideas of fragmented epistemology. Furthermore, I go over the various motivations discussed in the literature to opt for a fragmented epistemology, including fragmented models of belief. The section concludes by explaining how fragmentation delineates from contextualism, as well as with an overview over recent developments and open questions in the literature surrounding fragmentation.

2.3.1 The Cherniak model of fragmentation

In his book "Minimal Rationality", Cherniak (1986) argues against very idealised theories of rationality and develops a theory of minimal rationality which is accompanied by a fragmented picture of memory. The main driver for developing this theory is that he thinks orthodox theories of rationality are based on a highly idealised model of memory that imposes unrealistic demands on limited and fallible agents like humans (cf. Kindermann and Onofri, 2021).

He draws a distinction between short term memory (active memory) and long term memory (passive memory). He argues that short term memory has limited storage capacity and limited duration. But, unlike inactive long term memory, the agent can operate on the contents of short term memory. "Items that are active in short term memory can serve as premises for inferences [and] act as an input to practical reasoning"(Kindermann and Onofri, 2021). On the other hand, long term memory does not have to deal with the very limited capacity of short term memory. According to Cherniak, long term memory has "practically no capacity limit". But there

is a price to pay for this – the agent cannot directly and readily access what is stored in long term memory, and therefore, the agent cannot operate directly on the contents of information stored in long term memory. To be able to do so, the agent has to activate the memories and "move" them from passive long term memory into active short term memory.

Moreover, according to Cherniak, to optimise the process of finding and retrieving information from long term memory and activating it to make it available for reasoning and action, long term memory is compartmentalised / fragmented. Pieces of information in our long term memory are organised into subsets / fragments according to their subject matter / contents. This will make it easier for agents to effectively and quickly find pieces of information within the large amounts of information stored in long term memory. "Depending on the cognitive problem at hand, the agent will search within subsets that are most likely to be relevant to the problem in question" (Kindermann and Onofri, 2021). Through this fragmentation of long term memory, the efficiency and reliability of the search process can be optimised, since the agent does not have to randomly search through the depths of long term memory, but can search with the help of subject matter and "key words". The fragmentation in our long term memory allows us to conduct "fairly reliable searches within a limited about of time", but it can also lead to mistakes.

The main point of Cherniak is that mental fragmentation is central to a more realistic account of memory which is, in turn, central to his account of minimal rationality.

2.3.2 The Stalnaker / Lewis model of fragmentation

Lewis and Stalnaker have developed the account of fragmented belief as we know it today within the epistemological literature, by using the traditional possible worlds framework for propositions and belief, and introduce a fragmented version thereof.

To recall their account, I will, first, give a brief introduction into the concept of a proposition, as well how that concept is employed within formal

models of belief.

Propositions and possible worlds

Propositions are a central concept in epistemology, as well as formal and less formal approaches of belief and knowledge. Today, the standard conception of a proposition is in terms of possible worlds (Lewis, 1982, Stalnaker, 1984, cf. also Hoffmann, 2012). A proposition P, according to this standard conception, is a set of possible worlds; i.e. those worlds in which P is true[14]. Furthermore, a proposition P is true at a world w iff $w \in P$. Intuitively, one can think of a proposition as the semantic content of a thought, or utterance or assertion. Importantly, propositions are standardly taken to be the objects of belief and other attitudes, such as knowledge or doubt, as they are often thought to be the primary bearers of truth and falsity.

Since this conception of propositions depends on another notion, that of a possible world, one may legitimately ask what possible worlds are. Within this standard framework, the (logical) universe is usually taken to be a (finite or infinite) set of possible worlds W. Those worlds can be understood as different ways of settling possibilities, or different ways the world may be. Within contemporary (formal) epistemology and especially, within epistemic and doxastic logics, possible worlds are not taken to carry substantive metaphysical commitment, but instead are used as a tool of capturing concepts such as knowledge and belief. This does not mean that the metaphysical questions surrounding possible worlds are not interesting[15], but they are not the focal interest of formal epistemologists, epistemic logicians, as well as within this dissertation.

The conception of propositions in terms of possible worlds is by no means the only one. Alternative and competing conceptions of propositions include, most notably, the so-called structured account of propositions according to which propositions are complex entities that are structured and

[14] A slightly alternative way of formally modelling propositions is via functions from possible worlds to truth values, instead of sets of propositions.

[15] For the interested reader, an introduction and overview over metaphysical questions can be found in Menzel (2024).

have constituents that are bound together in a particular way (eg. Bacon, 2023, King, 2024). However, since the possible worlds conception and framework can be called orthodoxy in epistemology and since it is the basis of the version of fragmentation developed by Lewis and Stalnaker that this dissertation builds upon, further elaborating on alternative conceptions of propositions is beyond the scope of this dissertation.

Formal models of belief

I will now explain how this possible worlds conception of propositions is used in formal epistemology and epistemic and doxastic logic. This is important, since this framework is the fundament for the fragmented model of belief that Stalnaker and Lewis develop.

The semantics of epistemic (logic of knowledge) and doxastic logic (logic of belief) are most commonly given via possible worlds semantics. This means that, for example for a logic of knowledge, the truth conditions of formulae involving the knowledge operator K, applying to a proposition P are given via truth of P within possible worlds. More precisely, usually, knowledge of a proposition P is semantically defined via whether P is true in all possible worlds that are epistemically possible / accessible to the agent in question.

Since this dissertation, for reasons explained in the previous chapter, is not primarily concerned with the concept of knowledge but rather with belief and, even more importantly, justification (of belief), let me delve a bit into logics of justification and (justified) belief.

Doxastic logics or logics of belief aim to formally capture the concept and the logic behind belief. They study what kind of rules and theorems hold, given some initial and minimal assumptions on the concept of belief. These "minimal" assumptions often postulate belief to be consistent and closed under logic, this means that these assumptions are exactly what the orthodox model of belief presented earlier imposes. Based on these assumptions, the literature comes up with logical systems studying belief, most often using a possible worlds framework, where propositions are

captured as sets of possible worlds (cf. among others, Von Wright, 1951, K. Hintikka, 1962, K. J. Hintikka, 1969, Stalnaker, 2006, Van Benthem, 2006, Bjorndahl and Özgün, 2020, Rendsvig, Symons, and Wang, 2024). Related literature paths are logics studying not just belief, but justified belief, evidence-based belief and the like (cf., among others, Van Benthem and Pacuit, 2011, Baltag, Renne, and Smets, 2012, Baltag, Renne, and Smets, 2014), again, predominantly within possible worlds frameworks[16]. Furthermore, there are paths in the literature working on justification logics, again, often against the backdrop of a possible worlds semantics (cf. eg. Artemov, 2008, Artemov, 2019, Artemov, Fitting, and Studer, 2024). Justification logics, as well as AGM belief revision (Alchourrón, Gärdenfors, and Makinson, 1985, Grove, 1988) are also noteworthy in the context of this dissertation, because they employ so called plausibility models, that is, rankings of possible worlds by way of their plausibility. While I will explain what exactly this means in the next two chapters of the dissertation, it is important to keep in mind the inspiration for the employment of this formal apparatus[17].

To conclude this, many formal models of belief operate in a possible worlds framework, where propositions are taken to be sets of possible worlds. And, furthermore, often, propositions are taken to be the objects of belief. And belief has to be consistent and closed under logic[18].

Stalnakerian / Lewisian fragmentation

Lewis and Stalnaker are, among others, responsible for the development of possible worlds models of belief, as well as their establishments as orthodoxy within the philosophical community.

[16]There is also a large part of the literature of belief logics working on dynamic versions of belief logics, so-called dynamic epistemic logics, cf., among others, Van Benthem (1996), Gärdenfors (1988), Leitgeb and Segerberg (2007), Van Benthem (2007), Van Ditmarsch, Van der Hoek, and Kooi (2007).

[17]There is one more source of inspiration, not directly related to epistemic and doxastic logics, namely the literature on non-monotonic reasoning, cf. eg. Kraus, Lehmann, and Magidor (1990), Lehmann and Magidor (1992).

[18]Another very prominent way to formally model belief is Bayesianism. I will delve into Bayesianism in section 4.

Then, also Lewis and Stalnaker, develop fragmentation as a way to cope with the problem of logical omniscience that arises "for possible world accounts of propositional attitudes and their contents" (Kindermann and Onofri, 2021).

Stalnaker provides his own version of the unity principle presented earlier, rephrased explicitly using propositions as the objects of belief[19]:

1. If P is a member of a set of [believed] propositions and P entails Q, then Q is a member of that set.

2. If P and Q are each members of a set of [believed] propositions, then $P \wedge Q$ is a member of that set.

3. If P is a member of a set of [believed] propositions, then $\neg P$ is not a member of that set.

Stalnaker takes the agent's doxastic state to be represented by the set of worlds that represents their doxastic commitments, which are supposed to be; to obey these three rules.

Stalnaker defends a pragmatic-dispositional account of propositional attitudes. This means that he understands a propositional attitude such as belief to be primarily defined in pragmatic terms. According to this conception, we believe those propositions that we are disposed to act upon. To put it in Stalnaker (1984)'s words: "Beliefs, according to the pragmatic picture, are conditional dispositions to act. A rational agent is, in general and by definition, disposed to act appropriately, where what is appropriate is defined relative to his beliefs and desires. To say that an agent believes that P is to say something like this: The actions that are appropriate for that agent – those he is disposed to perform – are those that will tend to serve his interests and desires in situations in which P is true" (p. 82).

This also explains the strength of the unity principle / the three Stalnakerian rules for belief just presented. They need to be so strong to prevent irrational practical action. Those rules are (partially) there to ensure practical rationality. Since, imagine a belief set containing P and $\neg P$,

[19]The exact formulation is taken from Kindermann and Onofri (2021).

and a practical decision depending on whether P or $\neg P$ is the case. For example, imagine that you believe both that it will rain tomorrow and that it will not rain at all tomorrow. You want to decide whether to take an umbrella to work tomorrow or not. If it will rain, you take it, but if it will not rain at all you will not take an umbrella. But if you believe both that it will rain and it will not rain at all, you would have to both take and not take the umbrella, which is not only practically irrational, but practically impossible. You cannot both take and not take the umbrella with you.

But, now, Lewis presents an example that shows how our beliefs can be fragmented, and thereby, do not obey to the unity principle: "I used to think that Nassau Street ran roughly east-west; that the railroad nearby ran roughly north-south, and that the two were roughly parallel" (Lewis, 1982, p. 436). This set of beliefs is inconsistent and violates orthodox belief norms. Instead of completely getting rid of rationality postulates on belief, Lewis explains his set of beliefs using fragmentation: "Now, what about the blatantly inconsistent conjunction of these three sentences? I say that it was not true according to my beliefs. My system was broken into (overlapping) fragments. Different fragments came to action in different situations situations, and the whole system of beliefs never manifested itself all at once. The first and second sentences in the inconsistent triple belonged to – were true according to – different fragments; the third belonged to both. The inconsistent conjunction of all thee did not belong to, was in no way implied by, and was not true according to, any one fragment. That is why it was not true according to my system of beliefs taken as a whole. Once the fragmentation was healed, straight away my beliefs changed; now I think that Nassau Street and the railroad run roughly northeast-southwest" (Lewis, 1982, p. 436).

According to fragmentation, we can be in more than one belief state at once. In one context or circumstance, we could be in one belief state. And in a different context or circumstance, we could be in a different belief state that might be incompatible with the former one. In Stalnaker's words, "[a] person may be disposed, in one kind of context, or with respect to one kind of action, to behave in ways that are correctly explained by one belief state,

and at the same time be disposed in another kind of context or with respect to another kind of action to behave in ways that would be explained by a different belief state ... the agent might, at the same time, be in two stable belief states, be in two different dispositional states which are displayed in in different kinds of situations".

According to fragmentation, the three Stalnakerian rules have to be satisfied in each individual dispositional belief state / fragment, but across belief states, they do not have to hold. An agent's context-encroaching overall belief state may contain inconsistencies and does not have to be closed under logic. It is important to note that the unattractive consequences this can entail for practical rationality are avoided here, by claiming that it is only ever one single belief state (among the many that a fragmented doxastic state may contain) that is relevant for a particular circumstance or decision problem. Therefore, the violations of the rules of rationality will not ever influence the rationality of action, which is what having beliefs "is all about", according to Stalnaker.

Furthermore, according to both Lewis and Stalnaker, fragmentation is not something that is necessarily to be desired from a rationality standpoint, but something that happens in real life and that epistemology should be able to deal with. An agent who is fragmented is not necessarily irrational, as long as they can ensure practical rationality. But, as Lewis puts it, the fragmentation should ideally be "healed", because it is usually due some type of ill-justified belief, or mistake, like in Lewis case, a wrong way of approximating the directions of Nassau street and the railroad.

In this dissertation, I will diverge from this picture, and thereby untangle epistemic from practical rationality. Within my account of fragmentation, fragments are not necessarily tied to specific actions or practical decision making circumstances. It can, therefore, occur that two or more different and incompatible fragments might become relevant to a single specific practical decision problem. Such situations make it necessary to discuss defragmentation or the aggregation of fragments, which is something novel within the literature on fragmentation.

Furthermore, I will make the stronger claim that fragmentation, some-

times, is not just a "necessary evil", to result in a theory that scores higher on descriptive accuracy than epistemic orthodoxy does, but rather, the epistemically superior thing to do – the rational thing to do. I will argue that fragmentation can be purely epistemically rational in some situations, even if it might not be practically rational at the same time. Even if some doxastic state might not be able to provide rational practical action, it can still be epistemically rational. And there are reasons why one might not want to alter or get rid of such a state even if it poses problems for practical rationality. I will delve into what these reasons are in the dissertation.

2.3.3 The central theses summarised

Similarly to how I did for the orthodox picture, I will now provide both Kindermann and Onofri (2021)'s summary of the central theses of fragmentation, as well as the slightly more detailed summary of fragmentation provided by Bendaña and Mandelbaum (2021).

Kindermann and Onofri, 2021 give the following theses as a high-level synthesis of fragmentation:

1. The total set of an agent's beliefs (at time t) is fragmented into separate belief states.

2. Each belief state (at t) is a fragment whose constituent beliefs are consistent with each other and closed under logical consequence.

3. The belief states of a single agent (at t) are logically independent: They may not be consistent with each other, and the agent may not believe the consequences of his belief fragments taken together.

4. Different belief fragments of a single agent (at t) guide the agent's actions in different contexts or circumstances.

The main point where my account of fragmented justification will diverge from these tenets is tenet 4; different fragments might have to guide the agent's actions in the same context or circumstance. This also means that

the sources of and reasons for fragmentation will be different in my account; I will provide a much more detailed analysis about what can be reasons for (rational) fragmentation, a thing which is not to this degree and form existent in the literature on fragmented belief yet.

Now, let us take a look at the slightly more detailed version of fragmentation presented by Bendaña and Mandelbaum (2021):

1. Isolation: [Beliefs] [...] are stored in distinct, independently accessible data structures, which we call "fragments".

2. Inconsistency: Any fragment that harbours a belief that P and a belief that $\neg P$ will incite a revision process to eliminate the inconsistency, but having one fragment contain P and a separate fragment contain $\neg P$ need not.

3. Locality: Information updating, including belief revision, normally takes place within a single fragment at a time.

4. Redundancy: Different tokens of any particular belief may be stored in different fragments.

5. Multiple Resistance: Beliefs that are most resistant to revision are beliefs that are most redundantly represented.

The third, fourth and fifth thesis are not really present in Kindermann and Onofri (2021)'s summary presented above. The third thesis describes how fragmentation should work not synchronically but diachronically. It stipulates that belief revision and information updating is highly local; it happens within single fragments at a time.

According to the fourth and fifth theses, and contrary to what orthodox epistemology prescribes, it is fine for multiple belief tokens of the same belief type to be present in the agent's overall doxastic state. And the fifth thesis is a claim about how redundancy and belief revision interact. It tells us that the more token representations of a particular belief are part of the agent's overall doxastic state, the more difficult it is to get rid of

them via the process of belief revision (especially since belief revision takes place primarily locally, cf. tenet 3). Within my own account of fragmented justification, redundancy and multiple resistance will not play a big role; I focus on isolation and inconsistency, as well as closure under logic.

I will also not primarily focus on information revision within my account of fragmented justification. The dynamics of information change in fragmentation settings is very interesting, however, as already indicated, not the focal point of this dissertation, because it would go beyond its scope. Chapters 5 and 6, where I talk about fragmented Bayesianism, will take up parts of this topic; and the conclusion indicates that this would be a fruitful avenue for further research.

2.3.4 Motivations

I will now provide an overview over the motivations and evidence for fragmented accounts of belief that have been presented in the literature on fragmentation, both empirical and philosophical.

Logical omniscience and closure

I already presented how the orthodox model of belief and its unity thesis pose problems regarding logical omniscience. Let's briefly recall: Real agents do not adhere to the demands of the unity principle; their beliefs are often not consistent and often not closed under logic (cf. Parikh, 1987). Hence, unity is not descriptively adequate. But, it is not normatively adequate either. Normatively understood, the tenets of unity are not necessarily the ideal guide for the rationality of real agents, because real agents who fall short of these norms are not necessarily irrational. Harman (1986), for example, famously argues that "rational agents with limited storage capacity should not strive to draw any and all logical consequences form the beliefs they hold, or else their finite minds will become cluttered with trivial beliefs tat are irrelevant to their lives". In this dissertation, I will argue that there are other reasons beyond optimising cognitive capacity why not adhering to the norms imposed by unity and the orthodox model of belief

can be the rational thing to do, contrary to what orthodox epistemology imposes.

The first option that might come to mind how to weaken unity is to drop the constraints it imposes altogether. But this would result in an account of rationality (if the resulting account is normatively understood) that would be much too weak; pretty much anything would go in this case. The belief sets of rational agents must adhere to some normative standards, otherwise those agents cannot be called rational. And this would be a problem, because, to put it in the words of Cherniak, "the elements of a mind – and, in particular, a cognitive system – must fit together or cohere, a collection of mynath bird utterances or snippets of the *New York Times* are chaos, and so (at most) just a sentence set, not a belief set ... no rationality, no agent" (Cherniak, 1986, p. 6).

Among other options for weakening rationality constraints while not fully giving up on them described in section two, fragmentation accounts of belief have become one prominent way of dealing with this problem.

Fragmentation accounts of belief allow agents' beliefs to, overall, be not closed under logic, and to be, overall, inconsistent. But fragmentation does not drop rationality constraints altogether. Within fragments, the tenets of unity still holds, each fragment needs to remain consistent and closed under logic. But across fragments, inconsistencies are allowed and closure under logic does not have to hold.

Whether fragmentation can indeed be rational is one of the most discussed open questions within the literature on fragmentation. This thesis will expand on the already existing literature on rationality of fragmentation, by providing reasons for why fragmented justification be rational, as well as by giving a rationality argument for fragmented belief in addition to the arguments already present in the literature.

Explanatory power

Because the unity principle is descriptively inaccurate, it is not a good tool for predicting the reasoning and actions of agents (cf. Kindermann and

Onofri, 2021). It does not describe and therefore cannot correctly predict the reasoning and actions of real agents. Again, one might argue that this is not the goal of unity. But, if one takes this to be one of the goals of a model of belief, then the orthodox model of belief falls short of achieving this goal.

"If one wants a theory of belief that allows for the explanation and the prediction of the actions of real agents, we'd better let go of the extreme idealisations involved in unity" (Kindermann and Onofri, 2021). But, again, we do not want a fully descriptively adequate theory either; we do not want to get rid of all requirements on rationality. Fragmentation accounts of belief manage to be much more descriptively adequate, and thereby are better at predicting and explaining the reasoning and actions of agents, while not letting go of the normative dimension altogether.

Fragmentation accounts perform a balancing act between descriptive and normative adequacy – they aim to be descriptively more adequate than the orthodox model of belief while remaining normatively significant, albeit normatively weaker than the orthodox model of belief. Thereby, they manage to be better at explaining and predicting the behaviour of real agents than the orthodox model.

Evidence for fragmentation

I will now provide a (non-exhaustive) overview over some types of empirical evidence there are that favour fragmentation accounts of cognitive architecture and belief. This overview is largely taken from Bendaña and Mandelbaum (2021).

Extinction and reinstatement Consider a dog that has learned to associate the ringing of a bell (conditioned stimulus, CS) with receiving a treat (unconditioned stimulus, US). Now, through an extinction experiment, the learned association between the ringing of the bell and the treat gets broken. Imagine that the dog learned the association between the sound of the bell and the treat within his home. Now, the owner of the dog

takes it to a dog training school, i.e. a different context, where a trainer extinguishes the association by repeatedly ringing the bell without giving the dog a treat. The surprising thing is that, extinction displays failure. The dog comes back home and displays the seemingly unlearned association between the sound of the bell and food, eg. by salivating after hearing the sound.

Such extinction and renewal experiments are robust (Rauhut, Thomas, and Ayres, 2001, Bouton, 2004, Bustamante, Uengoer, and Lachnit, 2016). "Change of context can cause an immediate and robust return of a previously extinguished association" (Bendaña and Mandelbaum, 2021, p. 82). In different contexts, different information is accessible, and learning new information or retracting information is not necessarily transferred across contexts. As Bendaña and Mandelbaum (2021) argue, "fragmentation offers a natural explanatory framework for these results, as it allows for multiple token-representations of associations to be housed in different fragments, each activated by different contexts". The reoccurrence of a previously extinguished association happens, because there is a context-switch, where, although the fragment corresponding to the context where extinction has taken place does not contain token representations of the associations anymore, the extinction has not taken place in the fragment corresponding to the initial context. The orthodox model of belief cannot accommodate these kinds of patterns of information access, since it does not allow for multiple token associations of the same belief to be stored within the belief set.

Implicit bias Implicit bias experiments support a fragmented account of belief for reasons similar to the ones just presented in the context of extinction and renewal experiments. Implicit biases are inclinations or prejudices about people or groups of people that are often not believed explicitly but implicitly, which makes them more difficult to be aware of as the holder of the belief. This also makes them more difficult to get rid of. In the literature, they are often described as conditioned associations (De Houwer, 2009, Mandelbaum, 2015), similar to those described in the extinction and

reinstatement paragraph just discussed. Experiments aiming to get rid of or modulate peoples' implicit biases often struggle with the following phenomenon: After successfully eliminating the bias in a certain context, it often reappears upon retests for the bias after brief delays (Devine et al., 2012, Lai et al., 2016).

Again, fragmentation accounts of belief, unlike the orthodox model, can accommodate for these phenomena for similar reasons to the ones just explained regarding extinction and reinstatement.

Inconsistency & ballistic believing People constantly have inconsistent belief sets. Even really, really smart people like scientific experts, as I will show in chapter 4 of this dissertation, and authors of books, as I will show in chapter 3, as well as philosophers, such as Lewis (1982). (Seemingly) rational inconsistency does not fit at all into the orthodox model of belief; it does not allow for it, since it treats it as a threat to rationality.

Fragmentation accounts, on the other hand, do allow for inconsistent information to be stored within the agent's belief state. Within each fragment, orthodox rationality norms ought to hold, but across fragments, inconsistent information can be admissibly stored.

Psychology provides a lot of evidence in favour of agents having inconsistent sets of beliefs. One example of this is the theory of ballistic believing. This phenomenon, also referred to as the so called "Spinozan" model of belief acquisition, developed by Gilbert, Tafarodi, and Malone (1993), shows that, descriptively speaking, the process of acquiring beliefs is largely automatic and without filters for truth and falsity. The process of rejecting propositions, however, of deleting them from one's beliefs sets is non-automatic, but rather a "controlled, effortful and breakdown prone process" (Bendaña and Mandelbaum, 2021). We initially believe any truth-apt proposition we come across, and only later, and only when we have the resources, we sort out false ones by revising our belief state[20].

Since cognitive load is omnipresent, no one can be cautious enough to

[20]Another area of science supporting the same hypothesis is cognitive semantics, cf. eg. Evans and Green (2006), Fillmore (2006), Lakoff and Wehling (2012).

continuously reflect on all the propositions that enter one's belief state, especially those that one should reject because they conflict with previously accepted propositions and / or are false (Levy and Mandelbaum, 2014).

Fragmentation can accommodate this evidence, since it allows for inconsistent information to be stored within one's overall belief state, while the orthodox model is incompatible with this evidence. But, as already explained, the aim of the account of fragmentation developed in this dissertation will not be to adequately describe explain peoples' clearly irrational doxastic patterns, but instead, show how fragmentation, sometimes, can be the rationally superior thing to do.

Retraction Retraction experiments also support a fragmented account of belief instead of the orthodox model. Retraction experiments look like this: Participants of the experiment receive some information and then later are told that some of the previously presented information is false; some of the information is retracted. Then, the participants are being asked questions based on the information they were being presented with, including the retraction. Interestingly, participants still refer to the retracted information as being true although they have explicitly been told that that information is false. Such retraction experiments can be found, for example, in Wegner, Coulton, and Wenzlaff (1985), Johnson and Seifert (1994), and Seifert (2002)[21].

Unlike the orthodox model, fragmentation can explain these phenomena. Remember that one might argue that it is not the goal of the orthodox model to allow for / predict obvious failures of rationality, which these retraction experiments seem to be highlighting. This section is not about normativity, however, it is about empirical evidence supporting fragmentation accounts of belief over the orthodox model.

The widsom of the crowds The wisdom of the crowds effect describes the phenomenon that judgments made by crowds or groups are often more

[21] Other evidence supporting the same hypothesis can be found in Anderson and Kellam (1992), Sinnott-Armstrong (2008), or Mandelbaum and Ripley (2012).

accurate than the judgments of the individuals composing the group or crowd (Galton, 1907, Surowiecki, 2005)[22]. The average of the guesses / judgments of a group becomes more accurate, as the sources of error of those guesses become more independent from each other.

But how do the individuals constituting a group or crowd generate their individual guesses / judgments? Orthodox epistemology prescribes, but also it has been traditionally assumed by empirical researchers, that individuals are guided by the totality of the information stored within their doxastic state.

Interestingly, one can also detect the wisdom of the crowds effect for not groups, but within individual agents (Vul and Pashler, 2008). When letting individuals guess / judge multiple times on the same question, instead of just once, the error of the average of the guesses decreases significantly compared to the errors of each guesses, taken individually. Importantly, contrarily to what one might think, when people second-guess, there is no increase in the accuracy of the information available to them, nor does the total information available to the agent change.

The orthodox model of belief cannot deal with the wisdom of the crowds phenomenon within individuals. According to it, agents should have global access to all their stored information at once. This would prevent the possibility of the degrees of errors when individuals guess multiple times to be independent enough to produce the wisdom of the crowds effect. If the individual can globally access all information at once, first of all, the guesses should not change, but second of all, there should be no independence in the errors of the guesses.

Fragmentation accounts can explain the wisdom of the crowds effect in individuals. "[D]ifferent contexts [can] cause participants to access different fragments which contain different information" (Bendaña and Mandelbaum, 2021, p. 90). And consequently, guesses that come from information in one fragment have a different source of error than a guess produced by

[22]The wisdom of the crowds is also a much discussed topic in contemporary analytic philosophy, cf. eg. Mayo-Wilson, Zollman, and Danks (2013), Baccini et al. (2023), Dietrich and Spiekermann (2024) and Stewart and Elkin (forthcoming).

information from a different fragment. Allowing for multiple guesses concerning the same question allows agents to produce guesses accessing different fragments which makes it possible for the total accuracy to increase with second-guessing, having the independence of the sources of error increase.

These different pieces of evidence just presented show how the orthodox model is descriptively inaccurate – it cannot explain or predict these phenomena. Fragmentation, on the other hand, proves a useful framework for accommodating these phenomena. "People do not update globally. [And] their total set of beliefs is replete with redundancy and inconsistency" (Bendaña and Mandelbaum, 2021, p. 94).

Again, this discussion was not about whether the orthodox model or fragmentation accounts of belief are normatively adequate, but merely a presentation of empirical evidence in favour of fragmentation.

The agent's circumstances as fragment-determining

As already mentioned a few times, some pieces of information might be accessible to an agent in some circumstances, or relative to some task, but not in other circumstances, or relative to other tasks. Having information not be globally accessible but fragmented, with specific pieces of information only being accessible in specific circumstances, is one of the core motivations for fragmentation accounts, since they can accommodate this very well.

There can also be other epistemic and non-epistemic circumstances that determine the fragmentation process. One commonly mentioned one is subject matter. If fragments are activated by specific tasks or decisions, then it makes sense that the fragment being activated to make a decision on what to wear given the weather today will contain propositions on the weather, temperature, and appropriate clothing for certain temperatures and weather conditions, but not, say, information about classical music, politics, or philosophy. Leitgeb (2017) discusses further aspects

within the agents epistemic and non-epistemic circumstances that can be fragment-determining and context-inducing, such as what the agent cares about in some circumstances or others, or what her perspective is. Another fragment-individuating factor that will be important in this dissertation is that agents can sometimes rationally make different sets of background assumptions in different circumstances or contexts. And there will be more fragment-individuating factors that I will discuss throughout this dissertation.

One of the most important topics to discuss when it comes to fragmentation is whether it is merely an account that outperforms the orthodox model in being able to describe and explain the behaviour of real agents, or whether fragmentation accounts can also provide a meaningful notion of rationality. I will argue that fragmentation, sometimes, is the rationally superior thing to do, compared to the alternatives orthodox epistemology provides, and I will provide reasons for why fragmentation generally, and specific fragment-inducing factors, can be rational.

Solving epistemic puzzles

Fragmentation accounts of belief are attractive, furthermore, because they provide a solution to the preface paradox, which is one of the most important and widely discussed paradoxes within epistemology. The paradox is about the author of a book who seems to be justifiedly believing an overall inconsistent set of propositions. In the main part of the book, the author, who is a scientific expert on some specific topic, for example, baby nutrition, communicates her well researched and justified beliefs about baby nutrition, $A_1, ..., A_n$. But in the preface of the book, she admits that one of the many claims she makes in the main part of the book is false, i.e. $\neg(A_1 \wedge ... \wedge A_n)$. The orthodox model of belief, under a normative reading, would classify this author as irrational, because the author believes an overall inconsistent conjunction of propositions, $A_1 \wedge ... \wedge A_n \wedge \neg(A_1 \wedge ... \wedge A_n)$. But, intuitively, the author seems rational. Each claim she makes in the main part of the book about baby nutrition is well researched. But equally, the claim she

makes in the preface about there being a false proposition in the body of the book is well justified. The author, although being a scientific expert, is a human, and humans are fallible. Furthermore, she has published books in the past, and each time, mistakes have been discovered after the book has been published. This is the paradox – according to orthodox epistemology, the author is irrational in believing this overall inconsistent proposition of propositions, but it seems she is, in fact, rational in doing so, because each belief is well justified.

Fragmented belief accounts provide an attractive solution to this paradox that can explain why the author is not irrational while preventing the extreme consequences that completely breaking up with the orthodox model can bring about. The first one to develop a fragmented belief solution to the preface paradox is Cherniak (1986), more recently, Leitgeb (2017) has developed a detailed fragmented / contextual solution to the paradox (and, furthermore, the lottery paradox, too). Since I will explain how fragmented belief solutions solve the preface paradox in the next chapter, as well as how the account of fragmented justification I develop can provide an even better solution to this paradox, I will omit a detailed presentation of the solution at this point.

2.3.5 Delineation from contextualism

I have just mentioned that Leitgeb (2017) provides a contextualist / fragmentation solution to the preface paradox. However, the account Leitgeb develops is not explicitly a fragmentation account, but rather a contextualist account of belief. Furthermore, I mentioned how many authors within the fragmentation literature think that fragments "come" with a specific task or decision, and get activated in a specific circumstance or context. This begs the question what exactly the difference between a fragment and a context is, given that the two concepts seem very closely related. This is a question not yet investigated by the fragmentation literature that I will now proceed to answer.

I already distinguished two forms of contextualism earlier. There is tra-

ditional semantic contextualism, where the truth conditions of a concept change from context to context. And there is non-semantic contextualism, as developed and used by Leitgeb (2017), where the content of the concept in question does not change from context to context, but rather, the agent's epistemic and non-epistemic circumstances, such as her perspective, attention, or stakes etc. change with context. What can also change between contexts, here, is the types of content that are relevant to the agent. In one circumstance, information about the weather might be relevant, while in a different context, information about Gödel's incompleteness theorems might be more relevant to her. This is the type of contextualism I focus on here because it is closely related with fragmentation, unlike semantic contextualism.

But what is a context? And how does it differ from a fragment? One can perceive of the concept of a context in more or less mind-dependent ways. A context can – in principle – be perceived of simply as the (predominantly external / mind -independent) space-time situation an agent finds themselves in, including information about what time it is, where she is, who else is there, what the topic of reasoning or conversation is, and so on. But the way Leitgeb (2017) perceives of a context, and what I choose to adopt in this thesis, is that there is another mind-dependent dimension to a context; in different contexts, the agent's internal circumstances will be different, not just the "objective" circumstances of the external situation. What changes between contexts can be, as mentioned, the stakes, what the agent cares about, the information that is relevant, her perspective, among other aspects. These aspects can be perceived to have, both, internalist and externalist dimensions, i.e. a dimension purely internal to the agent, and a dimension that is external to the agent. For example, what the agent cares about is very internal to the agent, but the stakes or relevant information could potentially be imagined to be determined not just through internal reasoning of the agent, but through some external circumstances as well.

Thus, a context can be perceived as something purely internal and mind-dependent to an agent, but at least some external factors can play a role in determining the context, such as the physical situation the agent finds

herself in, the time, etc.

Fragments on the other hand are purely mind-dependent entities; they cannot be conceived of independently of the agent's mind, because they "live" within the mind. Fragments are one possible structure for how the agent's mental state can look like, as opposed to an unfragmented – global – mental state. In the fragmented belief literature, fragments are usually taken to be sets of propositions (or partitions over possible worlds). A context, on the other hand, is not a set of propositions. A context is what influences the agent's mental state. As the fragmentation literature perceives, different contexts correspond to different mental fragments. In a different context, with different internal and potentially external circumstances, different fragments will form and become active. The aspects of the contexts that the agent finds herself in in different situations influences how the different mental fragments will look like.

For contextualism, the idea is usually that, in different contexts, the mental state of the agent will look very different. Each context allows for a separate mental state and corresponding set of believed propositions. And within each context, the norms of the orthodox model of belief need to hold.

For fragmentation accounts, what contextualism allows for, i.e. having different mental states in different circumstances, is possible to be entertained by the agent, so to speak, simultaneously. Instead of having to commit to one global mental state, and having to completely switch that global state to a different one with changes in context, fragmentation allows the agent to introduce fragments; several sub-states of the agent's overall mental state that correspond to different contexts, and which will be active in different circumstances. And those fragments are allowed to make up the agent's overall mental state *simultaneously*; instead of having to discard her mental state and come up with a new one in a new context, the agent can dis-activate one fragment and simply move to another, without having to get rid of any information within her mental state.

One important topic of this dissertation will be de-fragmentation and the need for the aggregation of fragments in light of the demands of practi-

cal rationality which I will talk about shortly. This is another point where it is important to distinguish between a fragment and a context. Fragments are parts of the mental state of the agent and can be formally represented. For fragmented belief accounts, they are sets of propositions; in my account of fragmented justification, they will be formally represented via plausibility orderings of possible worlds, and, later in the dissertation, via credence functions. These are entities that make it possible to use social choice and judgment aggregation theory to study how one might get rid of having many fragments and merge them into one unfragmented mental state. Contexts are what fragments are influenced by, but they themselves cannot be aggregated or merged as easily, because they are more complex, and less easily explicitly formalisable.

2.3.6 Recent developments and open questions

Lastly in this part of the dissertation giving an overview over fragmented epistemology, I will provide an brief summary over recent developments and important open questions in the fragmentation literature.

First, some recent developments: Epistemological theories of belief or knowledge often posit more norms of principles beyond unity. Widely discussed principles within the literature on knowledge (but also belief) are iteration principles; i.e. principles that demand that if you know a proposition P, then you know that you know that P, and so on (cf. T. Williamson, 2000, Greco, 2015). Similar to the principles of unity – consistency and closure – one might argue against iteration principles like this on similar grounds (cf. Kindermann and Onofri, 2021 for a detailed example). Similar to consistency and closure, fragmentation allows for the relaxation of iteration principles. For example, one might know a proposition P relative to a certain fragment. But one might know the – different – proposition, KP, relative to a different fragment. In other words, iteration principles do not have to hold across fragments, but one can still preserve them intra-fragmentarily.

Fragmentation is also part of a relatively recent account of mathemat-

ical knowledge, developed by Rayo (2013). Rayo argues that "cognitive accomplishment in logic and mathematics can be modelled, in part, as the acquisition of information transfer abilities: Abilities whereby information that was available for the purpose of one set of tasks becomes available for the purpose of a different set of tasks". Fragmentation makes precise what the transfer of information amounts to: "A subject who has access to pieces of information for some purposes but not others is usefully thought of as being in a fragmented cognitive state. [...] One can model an information-transfer ability as the instantiation of a relation of accessibility amongst different fragments within the subject's doxastic state".

Furthermore, fragmentation has recently been studied in relation to and combined with models of belief that take belief to be question sensitive (eg. Leitgeb, 2017, Yalcin, 2018), and inquisitive semantics and logic (eg. Ciardelli, Groenendijk, and Roelofsen, 2018).

Recently, building on the original accounts of binary fragmented belief, that is all-or nothing belief, that have been developed by Lewis, Stalnaker & Co. (contrary to graded belief or degrees of belief, often modelled by Bayesian credence functions), fragmented versions of Bayesianism have received increased interest (Elga and Rayo, 2021, Elga and Rayo, 2022), but are still an underdeveloped topic that will receive attention in this dissertation, but should be studied beyond.

Open questions and topics that need to be discussed within the fragmentation literature are plenty. One of the most important questions concern what valid (and invalid) reasons for (rational) fragmentation are. My dissertation will add to this discussion a new perspective on what the reasons for fragmentation can be.

This is connected to potentially the most pressing question with regards to fragmentation – rationality. Fragmentation accounts are in a limbo state – wanting to be more descriptively accurate than orthodox models, but not wanting to completely lose their normative dimension.

In the early literature, fragmentation is often portrayed as a "necessary evil", something that happens to real agents, but something that clearly deviates from the ideal of perfect rationality (Stalnaker, 1984). Egan (2008)

and Egan, in his chapter in Borgoni, Kindermann, and Onofri (2021), argues that "fragmentation can serve as a useful damage-control device in cases where agents have belief forming mechanisms that are liable to go wrong. The fact that agents have such fallible belief-forming mechanisms can make it more rational, in certain kinds of cases, to be fragmented rather than unified". To rephrase this, because real agents reason with error-prone belief forming processes, it can be "rational" to be fragmented, given the error-prone belief forming processes.

Cherniak (1986), motivated by similar reasons, develops a theory of minimal rationality, and fragmentation is a part of this theory of minimal rationality. Minimally rational agents may well often make mistakes, and not be perfectly rational, but the might nonetheless exhibit a kind of minimal rationality.

So, the consensus so far in the literature about rationality of fragmentation is that fragmentation is definitely a deviation from ideal rationality, and the rationality that a fragmented agent may possess is a kind of minimal rationality that agents can possess although / while they make mistakes, are fallible, and computationally limited. My thesis will diverge from this consensus. I will argue that sometimes, fragmentation is the rational thing to do, and in those cases, fragmentation epistemically outperforms its more orthodox epistemic alternatives.

Moreover, there is some debate on inter-fragmentary rationality criteria; recall that the tenets of unity, within fragmented accounts, have to hold within but not across fragments. Borgoni (2021) and Yalcin (2021), both, discuss, whether, and if yes, which rationality norms should exist in between fragments. Because, although fragmentation should be less normatively demanding than the orthodox model, one might want to pose some limits on how many fragments there may be, how they should relate to each other and so forth. In my dissertation, I will discuss this topic and expand on the literature. I will provide what I call precautionary norms of fragmentation. And, furthermore, the account of fragmented justification I develop will be able to serve as an additional rationality argument for fragmented belief accounts, which is another benefit of my account.

Other open questions and topics in the literature I will not focus on in this dissertation include the relationship between fragmentation and language, fragmentation and mental files, as well as how fragmentation and the literature on implicit attitudes relate[23].

2.4 Social choice and judgment aggregation

The next topic that needs introduction is social choice, judgment aggregation, and opinion pooling.

Social choice theory has been pioneered by Condorcet, Borda and Arrow, in the 18th and following centuries, and deals with collective decision problems, and how to (best) arrive at group opinions and decisions given the opinions provided by the individual members of the group. Orthodox social choice theory deals with the aggregation of preferences (List, 2022), while the modern judgment aggregation literature deals with the aggregation of propositions / judgments (cf. List, 2012), and (probabilistic) opinion pooling is concerned with the aggregation of credence functions (Dietrich and List, 2016).

Condorcet initialised this field of study with, first, his jury theorem, and secondly, his famous paradox. Condorcet's jury theorem tells us that "if each member of a jury has an equal and independent chance better than random, but worse than perfect, of making a correct judgment on whether the defendant is guilty, the majority of jurors is more likely to be correct than each individual juror, and the probability of a correct majority judgment approaches 1 as the jury size increases" (List, 2022).

Moreover, Condorcet's paradox states that the preferences of the majority of a group can be irrational even if all the individual preferences are rational. Arrow, then, puts forward a more general account of studying preference aggregation; where "he consider[s] a class of possible [prefer-

[23]It is not being studied yet, but it would be fruitful to, furthermore, investigate fragmentation in the context of (formal) accounts of attention and awareness, cf. eg. Belardinelli and Schipper (2024). Related to this, I will discuss in chapter 5 how fragmented Bayesianism can be used to model different levels of attention.

ence] aggregation methods, which he calle[s] social welfare functions, and asked which of them satisfy certain axioms and desiderata" (List, 2022). He proves that there is no way of aggregation preference ranking into a collective preference ranking, of two or more individuals over three or more alternatives, where the aggregation rule satisfies a number of rationality conditions in the form of axioms[24].

This is a very significant result; since it tells us (and this translates to the aggregation of propositions and credences) that given a number of minimally rational requirements / axioms on the aggregation rule, the consistency and thereby rationality of the collective cannot be guaranteed. This means, in turn, that to ensure collective rationality, aggregation rules need to drop some of the seemingly very plausible rationality conditions / axioms. And, it will depend on the specifics of the aggregation problem which aggregation rule will best achieve the goals for that specific situation.

Now, how does this literature relate to fragmentation? Sometimes, as it will become clear throughout this dissertation, although it might be epistemically rational to be fragmented, practical rationality will require an agent to get rid of their mental fragments. This is a topic not discussed much in the existing literature on fragmentation, because traditionally it is assumed that, in any specific practical circumstance, at most one fragment will be active at once. However, I will show in this dissertation that some decisions or circumstances will require more than one fragment to be active / involved. Similar to the aggregation literature, if all fragments agree on everything, it will be very easy base one's practical decision on those fragments, or the aggregate thereof. But if they agree, then one might

[24]In the formal framework of judgment aggregation, the axioms amount to the following: Universal domain specifies the admissible inputs for the aggregation rule. Collective rationality specifies the admissible output of the aggregation rule. Systematicity and anonymity constrain the way the outputs are generated from the inputs. Systematicity requires that the collective judgment on a proposition P only depends on the individual judgments on P, and not another proposition, say Q. And anonymity requires all individuals to have equal weight in the aggregation (cf. List, 2012). The original set of axioms put forward by Arrow are the following: Universal domain, ordering, the weak pareto principle, independence of irrelevant alternatives, and non-dictatorship. They roughly correspond to the former in the sense that they achieve similar goals.

question the legitimacy of going fragmented in the first place for that agent, as I will argue. If the fragments do not agree on everything, then it will be more difficult to base one's practical decision on those fragments. The agent, then, sometimes, has to de-fragment in a way that ideally gets rid of or somehow resolves the incompatible information the different fragments might contain. To investigate how this can be done, I apply the social choice and judgment aggregation apparatus to fragmented doxastic states. The apparatus per se is not novel, but the application to fragmented agents is. Here, what is being aggregated is not the opinions, beliefs or preferences of multiple different individuals to create a group opinion, belief or preference. But here, different fragments that all belong to the *same* individual agent will be aggregated because some practical scenarios will demand for de-fragmentation. Interestingly, however, not all scenarios will require de-fragmentation. The preface paradox that I will discuss in the following chapter will be such a case.

Not only is the application of social choice and judgment aggregation novel in the context of fragmented agents, but the process of de-fragmentation itself is a topic that is under-explored in the literature on fragmentation and will be discussed in this dissertation.

2.5 Bayesianism

As already mentioned in the first part of this chapter, in addition to qualitative theories of justification, there are also quantitative or probabilistic ones; of which Bayesianism is probably the most widely studied and famous one.

The central idea of Bayesian epistemology is to use the mathematical probability calculus to model agents' degrees of belief or credences, as opposed to all-or-nothing beliefs that are studied in orthodox epistemology. An example of an all-or-nothing belief could be: "Climate change exists", whereas an example of a degree-of belief, or a credence could be: "It is really likely (say, 90%) that climate change will continue to increase".

Deductive logic is often taken to be the mathematical tool to model full / rational belief. Full beliefs are the kind beliefs that an agent fully, completely endorses. The probability calculus, on the other hand, is used to model not fully endorsed beliefs – degrees of belief. An agent can have a degree of belief (often referred to as "credence") in a proposition that is only 5 or 10 percent, but also 50 percent, or, in the limit, full endorsement, i.e. 100 percent[25].

Bayesian epistemology "can be described as the attempt to use an intuitive but powerful tool – the probability calculus – for tackling long-standing problems in epistemology and philosophy of science" (Hartmann and Sprenger, 2010, p. 1f). It models degrees of belief or credences as mathematical probabilities, where the probabilities are interpreted in a subjective or epistemic sense, as opposed to objective (or combinatorial) chances of an event.

Hartmann and Sprenger (2010) describe Bayesian epistemology as consisting of four pillars; the axioms of probability, the Dutch book argument, the principal principle, as well as the conditionalisation principle. I will briefly explain what these pillars are, and then go on to elaborate on what role Bayesianism plays for the research on justification in general, and for the topic of this dissertation, fragmented justification, in specific.

Axioms of probability Bayesianism models degrees of belief or credence using probability theory in the sense that a degree of belief or credence function needs to satisfy the axioms of probability. Let W be a non-empty set of possible worlds. Then, $\mathcal{A}$ is a so-called powerset algebra on W, i.e. $\mathcal{A}$ is the powerset / set of all subsets of W. Informally, this means that $\mathcal{A}$ contains all propositions (possible events) in W. Then $c : \mathcal{A} \to [0, 1]$ is a credence function iff it satisfies the following:

1) $c(P) \geq 0$, for all P

[25]There is a lot of research on how full / all-or nothing / qualitative belief and degrees of belief / credences relate; cf. Leitgeb (2017) for a very influential account on how the two cohere with each other.

2) $c(W) = 1$

3) For incompatible (mutually exclusive) propositions P and Q: $c(P \cup Q) = c(P) + c(Q)$

Dutch book argument Ramsey (1931) came up with the idea to explicate the concept of rational or irrational credence by referring to the "standard economic conception of rationality – irrational [credences] would cost us money if we let them guide our actions" (Hartmann and Sprenger, 2010, p. 2). According to Ramsey, whether or not a credence is rational depends on whether we would be inclined to reject or accept bets based on our credences.

Ramsey argues that sports events where bets are very common are not the only occasion where betting is relevant in our lives. "[...] [A]ll our lives are in a sense betting. Whenever we go to the station we are betting that a train will really run, and if we had not a sufficient degree of belief in this, we should decline the bet and stay at home" (Ramsey, 1931, p. 85). Our credences are interpreted by Ramsey as dispositions to bet; and they, under this interpretation, directly translate to the betting odds that can be considered fair[26].

But how does rationality play into this picture? How do we know that the rational credences are those that correspond to fair betting odds? That is what the Dutch book theorem tells us. It says that if our credences violate the axioms of probability, then the betting odds that are implied by our credences cannot be fair, and that, based on these resulting unfair betting odds, it would be possible for our opponent in the bet to make a risk-free gain on us by offering us a set of bets based on our betting odds (a so-called Dutch book). And this is why our credences, then, could not have been rational in the first place[27].

Notice how this means that the notion of (epistemic) rationality of an

[26]There is an isomorphism, a structure-preserving one-to-one mapping, between credences and betting odds

[27]Proofs of the Dutch book theorem can be found in eg. Kemeny (1955), or Skyrms (1980).

agent is directly tied to whether the agent is practically rational in the sense of whether their betting odds can be turned against them or not.

Principal principle The next pillar of Bayesian epistemology is the principal principle (Lewis, 1980 / 2010). This principle tells us how to deal with events where objective chances can be established, such as combinatorial games (eg. Nim[28]), or also the throwing of a dice, as opposed to events where objective chance cannot be known or established.

Lewis states that if an agent knows the objective chance of a proposition P to be equal to c, then, as long as they have no "overruling information", then their subjective credence in P must be equal to the objective chance of P.

Together, the axioms of probability, the Dutch book theorem, and Lewis' principal principle establish the static / synchronic dimension of Bayesian epistemology.

Conditionalisation Conditionalisation is Bayesianism's fourth pillar, and deals with the dynamic / diachronic aspect of credences. Conditionalisation tells us how agents should update their credence functions if they receive new information. According to conditionalisation, the appropriate credence in a proposition P can be arrived at by taking the conditional probability of that proposition P, given the new piece of information, E ($c\,(P \mid E)$)[29].

Together, the axioms of probability, the Dutch book theorem, Lewis' principal principle, and the conditionalisation rule make up the cornerstones of Bayesian epistemology[30] .

[28]Two players take turns removing (or "nimming") objects from a distinct heap of objects. On each turn, a player must remove at least one object, and can remove any number of objects provided they all come from the same heap. The goal is to either take the last object or avoid taking the last object.

[29]Using the Bayes Theorem, this conditional probability can be reformulated which makes it easier to calculate in practice.

[30]The interested reader can find more information on Bayesianism and Bayesian epistemology in, eg. Bovens and Hartmann (2003), or Sprenger and Hartmann (2019).

2.5.1 What does this have to do with this dissertation?

But what does this have to do with the project of this dissertation, which is fragmented justification?

Firstly, as already explained, in addition to the qualitative accounts of justification in orthodox epistemology, there also are quantitative / value property accounts of justification of which Bayesianism is the most famous and widely used one. One might wonder, where in the four pillars of Bayesian epistemology, does justification play a role? How is Bayesianism a theory of justification?

Bayesianism provides the tools to come up with many different notions of justification; often referred to as confirmation measures (cf. eg. Fitelson, 1999). I will provide an overview over those different notions of justification within Bayesianism in chapter 5.

Now, similar to qualitative accounts of belief in epistemology, also for Bayesianism, there has been criticism that the account is highly normatively demanding. For example, the problem of logical omniscience is also a problem that Bayesianism faces. Similar to the fragmentation movement in qualitative epistemology, researchers have come up with Bayesian versions of fragmentation (cf. Elga and Rayo, 2021, Elga and Rayo, 2022, Fleisher, 2023). However, Bayesian accounts of fragmentation are a much less studied subject than qualitative versions of fragmented belief. In this dissertation, I will provide a very general framework of fragmented Bayesianism as a theory of fragmented justification. It will be much more general than existing accounts of fragmented Bayesianism. And I will talk about a list of parameters that orthodox Bayesianism exhibits that can, each, be chosen to be fragmented or not, and depending on one's choices on the fragmentation of these parameters, the resulting version of fragmented Bayesianism will be more or less intensely fragmented, i.e. more or less close to orthodox unfragmented Bayesianism. It is important that this dissertation, in addition to the qualitative account of fragmented justification developed in chapter 3 and 4, also provides a quantitative account of fragmented justi-

fication which can be found in chapter 5. I will also give an inductive risk argument for fragmented Bayesianism, and compare the Bayesian version of fragmented justification to the qualitative account I develop earlier in this dissertation.

2.6 Imprecise probabilities

This is the last topic that needs introduction. Unlike orthodox Bayesians who think that one can / should model an agent's belief state by way of one single credence function that assigns precise credences to each proposition the agent can entertain, imprecise probability theory argues that sometimes there are propositions to which agents cannot (or should not) assign precise credences. Proponents of imprecise probabilism argue that, sometimes (or even often), agents might not have enough evidence or the evidence might not be conclusive enough to assign precise credences to all propositions. For example, take the proposition that it will snow on the 1st of January 2027 in Munich. You may have a rough idea, but even given evidence about the weather on the 1st of January in past years, it seems difficult to assign a precise credence to this proposition. Still, it the evidence might be good enough h to say something about one's credal state towards that proposition, instead of suspending judgment completely. For example, you might know that snow is quite likely in January in Munich, so you might want to assign a credence above 0.5, but beyond that lower limit, you might not be sure what exact credence to assign.

The idea of imprecise Bayesianism is that it is rationally permissible that, in such cases, one's credence is not a precise number, but rather an imprecise range of credences between an upper and a lower bound. Versions of this idea have been developed by Keynes (1921), Levi (1974), Jeffrey (1983), Kaplan (1983), Kyburg (1983), Gärdenfors (1988), Joyce (2005), Levi (1974), and R. Bradley (2009).

Formally, this is achieved by way of allowing an agent's belief state to be represented not by one credence function, but instead by a set of multiple

credence functions. The interpretation of this formalism is – canonically – as follows: If all individual credence functions agree on what (exact) credence to assign to a proposition P, then the overall credence of the agent in that proposition is precise, it is exactly the credence that all individual credence functions agree on. But if the individual credence functions disagree over what credence to assign to a specific proposition (i.e. they assign different credence values to the same propositions), then the overall credence of the agent in that proposition is imprecise; and amounts to an interval of credences, with the lowest assigned credence by any individual credence function as the lower bound and the highest assigned credence by any individual credence function as the upper bound.

Within the imprecise probability literature, there are two camps with respect to the preferred answer to the question whether it can happen that rationality requires imprecision. Joyce (2005) argues that one's credences should never be sharper or more precise than one's evidence requires, and that if the evidence does not require or permit a precise credence, then assigning imprecise credence is rationally required (cf. Mahtani, 2019). The other camp argues that it might sometimes be rationally permissible to assign to have imprecise credences, but can never be positively rational, or rationally required.

The reason for introducing the topic of imprecise probabilities is that the formal framework of imprecise Bayesianism is extremely similar to the formal skeleton of accounts of fragmented Bayesianism. They also share motivations, as well as some related discussed topics in the respective literature streams. However, the literature streams do not relate to each other at all. Therefore, in chapter 6 of this dissertation, after developing my account of fragmented Bayesianism, I will compare the accounts of imprecise Bayesianism and fragmented Bayesianism, and comment on and evaluate their similarities and differences.

2.7 Additions to the literature

I will now provide an overview over what additions this dissertation will provide to the respective streams of literature just presented.

2.7.1 On justification

This dissertation provides a novel account of fragmented justification. The account will be an addition to the literature on justification, since existing orthodox accounts of justification are all such as evidentialism and reliabilism, are non-fragmented and are also implicitly based on a non-fragmented picture of rationality. The account of fragmented justification that I will develop in this thesis will be an attractive account to choose, should one commit to a fragmentation picture of belief, and of rationality more broadly.

I will exemplify how existing accounts of justification are based on a non- fragmented epistemology using evidentialism and reliabilism as two prominent case studies of traditional accounts of justification.

Evidentialism is often thought of as the default or common sense account of justification. As already explained, it requires one to have good reason for believing in a proposition to, in fact, be justified in believing it. There being many different options for precisely spelling out this account, a standard formulation of evidentialism can be found in Conee (2004): "Doxastic attitude D toward proposition P is epistemically justified for an agent S at time t if and only if having D toward P fits the evidence that S has at t." Evidentialism is most commonly conceived as a coherentist account of justification. Coherentism about justification is the view according to which justified belief is structured like a web, and the justification of belief should be understood in terms of the way in which the proposition believed coheres with other believed propositions.

Recall also that (process) reliabilism on the other hand makes the justification of a belief depend on the reliability of the process producing or causing it (see i.e. Goldman, 1979). According to Goldman (2012)'s process reliabilism, reliability "consists in the tendency of a process to produce

beliefs that are true rather than false". Reliabilism is usually conceived as a foundationalist theory. Foundationalism about justification is the view that justified beliefs / the justification of belief is structured like a building; they are divided into a foundation (consisting of basic beliefs which are not justified by any other beliefs) and a superstructure (consisting of non basic / inferentially justified beliefs). For reliabilism, this means that for a basic belief to be justified, it needs to result from a reliable belief-independent process. And a non-basic belief is justified if it results from a conditionally reliable belief-dependent process.

I will now argue for why any coherentist or foundationalist theory about justification, and thus evidentialism and reliabilism as examples from the two camps, is not based on a fragmented picture of rationality. For evidentialism as a coherentist theory, any (justified) beliefs ought to be a member of one singular all-encompassing web of beliefs which needs to be coherent, at least consistent. What justifies beliefs, here, is they way in which each belief relates to other beliefs in this coherent set. Whether a belief is justified or not is thus a global question – it can only be answered looking at its status in the total belief web. This leaves no room for fragmentation, because for fragmentation, the total belief set would need to be allowed to be split up in smaller belief sets which might contain cross-fragment inconsistencies. And then, the question about the justification status of belief would not necessarily be a global one. An implicit assumption that is made by coherentist theories of justification and evidentialism more specifically, is that it is impossible to have equally good reasons or evidence for a proposition P as well as its negation $\neg P$ at the same time[31].

[31] Permissivism is a sub-account of evidentialist theories of justification that, one might argue, is based on a fragmented picture of rationality, which would refute the argument. Permissivism is commonly referred to as the thesis that a source body of evidence can permit more than one rational doxastic attitude towards a certain proposition (cf. Jackson and Turnbull, 2024). Strong versions explicitly allow for the rationality of cases where an agent rationally believes a proposition P given a set of evidence, where it is consistent with the agent being fully rational and possessing the same set of evidence, and believing $\neg P$ instead (cf. White, 2005). There are two main reasons for why, although this may sound similar to fragmentation, permissivism is not a fragmented account of justification. First, for permissivists, justification or support is not fragmented. Although a certain body of evidence might support different and potentially opposing

Let me turn to foundationalist theories and reliabilism as an example of a foundationalist theory of justification. Recall that for foundationalism, all non basic justified beliefs are somehow derived, potentially via other non-basic justified beliefs, or from basic or foundational justified beliefs. It is not specified what exactly is intended by the term of a "derivation" – i.e. whether it should be interpreted in a strict sense of classical logic or whether a more loose interpretation would also be possible. However, even if a more loose interpretation is possible, a minimum requirement must be that the beliefs that are (directly or indirectly) derivable from each other must be consistent with each other. If they were inconsistent, it would be hard to imagine how they could be in principle derivable from each other.

But what about the basic beliefs that the derived beliefs are based on: Do they need to be consistent with each other? If one assumed there to be more than one basic belief and cross-derivation to be possible (i.e. that it is possible that a belief B rooting from basic belief P could be derived via a series of beliefs that originate in a different basic belief Q), then it makes sense to assume that all basic beliefs should be consistent with each other, too. In the special case of not multiple but just one basic belief, the question whether basic beliefs ought to be consistent would become redundant. Furthermore, if cross-derivation is possible, this then makes it

---

conclusions, this does not render justification per se fragment-relative. For permissivists, there is one big unified doxastic state, within which different attitudes towards the same proposition may be equally rationally permissibe. But all these attitudes would "live" within the same unfragmented unified doxastic state. Of course, the permissivist then needs to tell a story about how an agent can entertain equally well justified doxastic attitudes towards potentially inconsistent propositions within one doxastic state which is traditionally viewed to have to be consistent in order for the agent to qualify as rational. One attractive option to do so would be to adopt a weaker logic, or to, indeed, go fragmented, but permissivism as such does not assume any kind of fragmentation, neither of justification, nor belief or rationality. Fragmentation would be attractive to combine with permissivism, since it could tell a story about why it can be rational for an agent to have conflicting doxastic attitudes beyond merely saying that the body of evidence is "weak" enough to permit it. Secondly, often, permissivism is spelled out in a way such that the multiple and potentially conflicting doxastic attitudes that an agent may have towards a proposition given a body of evidence cannot be held at the same time. Most versions merely state that it can be equally rational to believe eg. either P or $\neg P$ given that body of evidence, but it would not be rational to believe both at the same time.

necessary that all basic and non-basic beliefs ought to be consistent with each other.

It results from these considerations that, as for coherentism, for foundationalism, the total hierarchy of basic and non-basic beliefs should be consistent. If inconsistencies are not possible, this is a strong argument against this account being based on any kind of fragmentation picture. Furthermore, again, the question whether a specific belief is justified or not, is a global question. It cannot be considered in a specific isolated context or fragment of the belief hierarchy, instead, the entire justification history of this belief might need to be considered, and is, at least, relevant to the justificatory status of the belief in question. As for coherentism, the traditional view of foundationalism is not fragmented.

However, this does not mean that either coherentism or foundationalism might, in principle, cannot be compatible with a fragmentation – i.e. that one could not come up with a fragmented version thereof. In fact, my account of fragmented justification will have both, a foundationalist, and a coherentist dimension. Here, I am making the weaker claim that the original versions of coherentism and foundationalism, and thereby evidentialism and reliabilism, are based on a non-fragmented picture of justification, belief, and rationality more generally.

One might argue, trying to refute my argument, that there are contextualist accounts of justification present in the literature and that they could be interpreted as versions of a fragmentation account of justification. And therefore, one might argue that fragmented accounts of justification, implicitly, have already been part of the literature. A contextualist account of justification was first presented by Annis (1978), other versions are by S. Cohen (1999) or Wedgwood (2008). As already explained, these accounts of contextualism about justification are traditional in the sense that they are semantic in nature. This means that the content and truth conditions of the term "justified belief" change from context to context. According to Annis' original account, the issue context of a claim is of central importance for determining whether or not they are justified in believing that claim, since it determines "the level of understanding and knowledge required" (p.

215) for justification. So, in some contexts stronger notions of justification hold than in others. The second constituent of his account is the ability of the agent to meet objections by an appropriate objector group made of critical truth-seekers. In this way, his account, besides being contextualist, also has a social dimension. For example, if the claim is about climate science, then the appropriate objector group might be climate scientists. Basically, this condition amounts to a "no defeaters" condition, according to which the claim should not be undermined by any counter-evidence the objector group might possess. S. Cohen (1999) and Wedgwood (2008) go in a similar direction, claiming that there are different standards for justification in different contexts.

While such a traditionally contextualist account of justification is compatible with fragmentation, this kind of contextualism should not be confused with fragmentation for the following reasons. The contextualism employed in those account is of a semantic kind. The content and truth conditions of the term "justification" change from context to context. This will not be the case for fragmentation about justification – here, the semantics of the term "justification" remain constant. What changes from fragment to fragment are the agent's epistemic and non-epistemic circumstances, but not the semantics of justification. Furthermore, fragments and contexts are not the same thing. Fragments are mind-dependent entities – fragments are ways for an agent to separate information, whereas the term "context" usually has both a mind-dependent as well as a mind-independent dimension (aspects of the context can also be "out there in the world", for example, a seminar room could (partially) constitute a context). For the reasons just presented, although there are contextualist accounts of justification already present in the literature, they are not fragmentation accounts of justification.

Given the vast and growing literature on fragmented belief and rationality, my account provides the advantage that it is a natural choice, should one want to commit to a fragmentation picture of belief and rationality. This is not to say that any unfragmented account of justification already present in the literature is per se incompatible with fragmented accounts of

belief and rationality, but that it might be desirable to add an account of fragmented justification to the literature on fragmented belief, knowledge and rationality.

2.7.2 On fragmentation

This thesis will add to the literature on fragmented epistemology a qualitative (non-probabilistic), as well as a quantitative (probabilistic) account of fragmented justification, in addition to the already existing accounts of fragmented belief and fragmented knowledge.

The following chapter 3 can be viewed as developing the skeleton and central idea of an account of fragmented justification. And it is a case study; it showcases how this novel account can provide a new and attractive solution to one of the central paradoxes of epistemology, the preface paradox (Makinson, 1965). Unlike existing solutions, this solution employs the central concept involved in the paradox – justification. First, I present the original paradox, and recall a prominent family of solutions to it – solutions that make use of fragmentation accounts of belief. I argue that, while fragmented belief solutions go in the right direction, what does the work in this paradox is not belief but justification. I then present my account of fragmented justification. I show how it solves the preface paradox in a novel way and I provide reasons for why it is a good solution. In particular, I will show how this account of fragmented justification and solution to the paradox outperforms other solutions, but also how it manages to provide a rationality argument for fragmented belief accounts and corresponding solutions to the paradox.

The upshot of my qualitative account of fragmented justification is that the justification relations that hold between pieces of evidence and the propositions they support are fragment-relative. And, fragments will be plausibility orderings over possible worlds.

Then, chapter 4 develops a more full fledged account of qualitative fragmented justification, using this time not using the preface paradox but a scenario from climate science as a motivating example for the account.

Here, a scientific expert about climate science seems to be justified in believing an overall inconsistent set of propositions. The chapter explains how the toolkit of orthodox epistemology is limited in dealing with situations like this, and argues that fragmentation is a more appropriate strategy than existing approaches to deal with (seemingly rational) inconsistent doxastic states. The chapter provides a formalisation of this example. The upshot of this formalisation is that justifications are represented via non-monotonic fragment-indexed relations that hold between propositions that need justification and evidence, where fragments are plausibility orderings over the universe of possible worlds, W. It expands on a number of justifying reasons for why fragmentation can be rational. It, also, develops some precautionary norms of fragmentation. Finally, it concludes by explaining how fragmented justification relates to belief, and argues that fragmented justification is a novel rationality argument for accounts of fragmented belief, and, furthermore, how one should go about de-fragmenting in light of practical rationality.

In the last two chapters, 5 and 6, I develop a quantitative version of fragmented justification – fragmented Bayesianism.

2.7.3 On Bayesianism

In this dissertation (chapter 5), I will develop a novel and general version of fragmented Bayesianism that goes beyond the accounts of fragmented Bayesianism that already exist (Elga and Rayo, 2021, Elga and Rayo, 2022, Fleisher, 2023). This account will showcase all the options there are within Bayesianism to go fragmented. It provides a list of parameters along which the account can be made moderately or very fragmented, depending on which and how many of those parameters one chooses to relativise to fragments. It, furthermore, presents an inductive risk argument in order to justify fragmentation. Lastly, I will compare the qualitative account of fragmented justification I develop in chapters 3 and 4 with the account of fragmented Bayesianism developed in chapter 5.

2.7.4 On social choice and judgment aggregation

In this dissertation, I evaluate how fragmented justification relates to belief, and practical rationality. As I will show, sometimes, but not always, will it be necessary, in light of practical rationality, to do something about one's fragmented mental state. I will present and use the framework of social choice and judgment aggregation as a tool to evaluate how to go about the process of de-fragmentation.

This is novel for two reasons; the first is that this framework has hitherto not been applied to the context of fragment-aggregation, and the second, related reason is that the topic of de-fragmentation has not been discussed much so far in the fragmentation literature. The reason for this is that, so far, fragments have been presented in a way that only at most one fragment will ever be relevant to a particular decision making problem, which prevents de-fragmentation from ever becoming necessary, since, internally, fragments adhere to traditional rationality norms and thus do not pose threats to practical rationality.

I will challenge this. The way I present fragmentation and its benefits, it can be the case that multiple fragments become relevant to a single decision – this will sometimes require de-fragmentation on pain of practical irrationality. I provide an analysis of how, if rationality requires, one may get rid of fragmentation. I will argue along the lines of the famous inductive risk argument by Rudner (1953) that the employment of non-epistemic values will be necessary in choosing an appropriate de-fragmentation strategy in any specific circumstance.

2.7.5 On imprecise probabilities

There is an obvious connection between fragmented accounts of Bayesianism and the imprecise probability framework – the underlying formal frameworks are (almost) identical. In chapter 6 of this dissertation, I will explore this hitherto unexplored connection between fragmentation and imprecision, and delve into the similarities and differences between the two frame-

works.

Chapter 3

The Preface Paradox and Fragmented Justification

This chapter presents a new solution to the preface paradox (Makinson, 1965). It makes use of a fragmented account of justification, and thus, unlike existing solutions, employs the central concept involved in the paradox – justification. First, I present the original paradox, and recall a prominent family of solutions to it – solutions that make use of fragmentation accounts of belief. I argue that, while fragmented belief solutions go in the right direction, what does the work in this paradox is not belief but justification. I then present my account of fragmented justification. I show how it solves the preface paradox in a novel way and I provide reasons for why it is a good solution. The upshot of my account of fragmented justification is that the justification relations that hold between pieces of evidence and the propositions they justify are fragment-relative.

Note that an article version of this chapter, at time submission to UB München, after the defense of this dissertation, has been accepted (forthcoming) for publication by Philosophical Studies – An International Journal for Philosophy in the Analytic Tradition.

3.1 The preface paradox

An author is about to publish a book on baby nutrition. She is an expert on the topic and has done a great amount of research for this book. On the basis of this research, she believes a large number of propositions A_1,...,A_n about baby nutrition and consequently publishes the sentences expressing these propositions in her book. But in the preface of the book, she admits that at least one of these propositions is false, thus that it is not the case that A_1,...,A_n are true. Therefore, she seems to be believing the following inconsistent conjunction of propositions: $A_1 \wedge A_2 \wedge ... \wedge A_n \wedge \neg(A_1 \wedge ... \wedge A_n)$. This "appears to present a living and everyday example of a situation [... where] it is [...] rational to hold logically incompatible beliefs" (Makinson, 1965, p. 205).

Makinson argues for the rationality of a situation like this in terms of the justifications of the author's beliefs: The author is justified in believing the propositions she expresses in the main part of the book because she has done a great amount of research to support her claims; the propositions she expresses in the book are well justified. But on the other hand, the author also has justification that leads her to believe that not all propositions in the book are true. The justification for this belief is that she has written many books before and has always received corrections from readers and reviewers; i.e. false propositions have been discovered after the books have been published every time. But, even if this were her first book, she would have justification for the preface proposition; namely, that she, as every human, is fallible.

This original version of the paradox is entirely qualitative. Note that Makinson does not talk about degrees of belief or credences. The paradox can be rephrased in terms of credences and a number of solutions apply to such reformulations of the paradox (cf. eg. Douven, 2002, Cevolani and Schurz, 2017, Kauss, 2021). Here, I will stick to the original qualitative version of the paradox.

Makinson solves the paradox by arguing that the author can rationally believe each proposition in the book, including the preface proposition. But

she cannot rationally believe the set containing all propositions inside the body of the book and the preface proposition at once: $\{A_1, ..., A_n, \neg(A_1 \wedge ... \wedge A_n)\}$.

Since its first formulation, many different solutions besides Makinson's original one have been proposed. Many of them involve weakening or denying closure under conjunction to block the problematic conjunction $A_1 \wedge A_2 \wedge ... \wedge A_n \wedge \neg(A_1 \wedge ... \wedge A_n)$ from arising. One prominent way of doing so is via fragmentation accounts of belief (cf. Cherniak, 1986, Clarke, 2017, Leitgeb, 2017). The new solution that will be presented here – fragmented justification – is related to fragmented belief solutions in that it is a fragmented account, too. However, both the account of fragmented justification and the solution it provides to the paradox will, firstly, exhibit significant differences to fragmented belief. Furthermore, the new solution presented here, unlike existing solutions, will be able to capture the fact that justification, and not primarily belief, is responsible for producing the paradox (specifically for the intuition that the author is rational). Lastly, the account of fragmented justification will provide a rationality argument for fragmented belief, and therefore increase the attractiveness of fragmented belief solutions to the paradox.

The chapter will be structured as follows: I will first recall fragmented belief solutions to the paradox. Then, I present the account of fragmented justification. In the following solution section, I show how the account solves the paradox. Finally, I argue that this solution is attractive, because, unlike other solutions, it solves the paradox using the central concept for the preface paradox – justification. A second argument for this account will be that it provides a rationality argument for fragmented belief, and therefore, indirectly, improves the attractiveness of fragmented belief solutions, which is an additional benefit of the solution developed here.

3.2 The fragmented belief solution

To later discuss and evaluate the differences between the solution developed here and fragmented belief solutions, and to discuss the advantages the account developed here has over fragmented belief, I now recall the central points on fragmented belief and how it solves the paradox.

Traditional accounts of belief demand that an agent's total set of beliefs should be consistent and closed under logic (cf. eg. Quine, 1951, Cherniak, 1983, Fodor, 1983). If an agent does not adhere to these norms, they are irrational.

This traditional picture is demanding. Real agents will almost never be able to adhere to these strong norms (Rauhut, Thomas, and Ayres, 2001, Bouton, 2004, Porot and Mandelbaum, 2020, Bendaña and Mandelbaum, 2021). This is not necessarily problematic, one might argue, because this normative picture does not have to be descriptively adequate. And moreover, it might be accurate that real agents are just not rational, fallible, make mistakes, and rarely believe in all the consequences of their beliefs. But, what does seem problematic about traditional accounts is this: It classifies as irrational even highly epistemically able agents which seem to, in fact, be rational, although they might fail to adhere to these traditional rationality norms. An example of such a situation is the preface paradox – the reason why the paradox arises is that although orthodox epistemology classifies the author as irrational, the author seems rational nonetheless.

In response to these traditional accounts, a number of philosophers have developed fragmentation accounts of belief, such as Davidson (1982b), Lewis (1982), Cherniak (1983), Stalnaker (1984), Cherniak (1986), Egan (2008), Rayo (2013), and Greco (2015). The central idea of such accounts is to weaken traditional rationality postulates on belief while remaining normatively significant enough to avoid rationalising agents making mistakes and thereby trivialising the notion of rationality. This is achieved by relativising belief to fragments.

Fragments are independently accessible proper subsets of the agent's to-

tal set of beliefs (Bendaña and Mandelbaum, 2021)[1]. Formally, for these accounts, beliefs are usually modelled as propositions (sets of possible worlds) and belief-propositions are taken to be fragment-indexed, i.e. they don't hold generally, but only relative to a specific fragment. What changes from fragment to fragment are the agent's epistemic and non-epistemic circumstances, such as the subject matter or perspective (cf. Leitgeb, 2017).

The rationality norms of fragmented belief accounts are such that each fragment, internally, has to adhere to traditional rationality norms, i.e. remain consistent and closed under logic. But across fragments, these traditional norms do not have to hold (Borgoni, Kindermann, and Onofri, 2021). For example, if an agent believes A in fragment F_1 and believes $\neg A$ in a different fragment F_2, this can be rationally permissible assuming that belief is fragmented.

Fragmentation accounts do not deny the rationality of an unfragmented and overall consistent belief state. But they add the thesis that a fragmented belief state that might contain inconsistencies across fragments can be rational, too[2]. Thereby, fragmented belief accounts claim that an agent need not live up to the strict rationality standards that the traditional picture of belief imposes. An important question regarding this account is whether a doxastic state fragmented in this way can indeed be rational. In this chapter, I will argue that the account of fragmented justification will provide a rationality argument for fragmented belief.

How does fragmented belief solve the preface paradox? Fragmented belief solutions argue that the propositions that the author communicates, respectively, in the preface and the body of the book, belong to two different fragments. They do so on the basis of the difference in subject matters of propositions in the preface and the body of the book. The preface is about the author's track record with regards to making mistakes when publishing books or general human fallibility, and the body is about baby nutrition. The author believes the preface proposition $\neg(A_1 \wedge ... \wedge A_n)$ rela-

[1]Sometimes, fragments are also modelled via partitions on W (cf. Leitgeb, 2017).

[2]For a discussion on the rationality of fragmented belief cf. Borgoni, 2021, Yalcin, 2021.

tive to the preface fragment, and the propositions $A_1, ..., A_n$ relative to the body fragment. Because closure under conjunction does not apply across fragments, the inconsistent conjunction of these propositions is blocked. The author remains rational believing each proposition involved relative to the appropriate fragment.

Fragmented belief solutions to the preface paradox are attractive – by relativising belief to fragments, they can explain how the author remains overall rational while believing an inconsistent set of propositions. But, there are two deficiencies of fragmented belief solutions. The account of fragmented justification will, as I will argue, in contrast, not face them.

First, recall that Makinson argues that the main reason why the author is rational in believing the propositions $A_1, ..., A_n$ as well as the preface proposition $\neg(A_1 \wedge ... \wedge A_n)$ is that every proposition is well justified. The first deficiency of fragmented belief solutions is that the importance of the justification for the propositions involved in the paradox is not adequately reflected in these solutions. The fragmented belief solution primarily makes use of the difference in subject matter between the preface and the body of the book to justify the introduction of fragmentation within the author's belief set.

This deficiency holds not only for fragmented belief solutions. The crucial role that justification plays in this paradox has not been discussed much in the literature on the preface paradox at all[3]. The account of fragmented justification will solve this deficiency of existing solutions, and specifically fragmented belief solutions, by providing a solution to the paradox that keeps the attractive features of fragmentation, but makes use of the concept of justification. I will argue that that what should be fragmented and does the main work in both rationalising and solving the paradox is not primarily belief, but justification.

[3]The only exceptions are the following: It is discussed indirectly within quantitative reformulations of the paradox in terms of coherence between belief and degrees of belief (cf. Moser and Tlumak, 1985). It is also discussed by Williams (2021) within a more qualitative context, who argues that rationality requires one to sometimes hold beliefs in jointly inconsistent but justified propositions. However, here, justification (in terms of evidence) only plays a peripheral role.

And second, an important question regarding the account of fragmented belief, and thereby about the solution it provides to the paradox, is whether being fragmented with respect to belief can indeed be rational. The solution hinges on the thesis that the author's fragmented belief state is rational. But this requires argument. The second deficiency of fragmented belief solutions is that it does not provide a strong argument for why fragmentation is rational in the context of the preface paradox.

The account of fragmented justification will solve this issue by providing a rationality argument for fragmented accounts of belief. This will strengthen fragmented belief solutions. Thus, besides providing an interesting solution by itself, the account of fragmented justification will be able to boost the credibility of fragmented belief solutions, which is another benefit of this account.

3.3 Fragmented justification

In this section, I present a sketch of an account of fragmented justification, focusing on the aspects relevant for solving the paradox. One may view this chapter as a case study that is the starting point for a fuller development of an account of fragmented justification.

Let W be a non-empty and finite[4] set of possible worlds[5]. Let $E_1, ..., E_n$ as well as $A_1, ..., A_n$ be propositions in the standard possible worlds sense (Stalnaker, 1978, Stalnaker, 1999); i.e. sets of possible worlds. A proposition, E, or A, is true at a world w just in case w is an element of E or A. $E_1, ..., E_n$ are pieces of evidence. $A_1, ..., A_n$ are the propositions in need of justification[6]. Fragments are modelled via total plausibility pre-orders $\leq_1, ..., \leq_n$ on the set of possible worlds W – they order the worlds in W

[4]For reasons of simplicity.

[5]I use a possible worlds setup in my account, since it is a standard practice in modelling propositions and varieties of doxastic attitudes.

[6]This distinction is terminological and for ease of exposition – it might well be that a proposition that plays the role of a piece of evidence E, itself is in need of justification, and therefore is an A-proposition, too

from most to least plausible[7].

Then, let $\mid\!\sim_{\leq_1}, ..., \mid\!\sim_{\leq_n}$, be fragment-indexed non-monotonic justification relations that connect one or many piece of evidence $E_1, ..., E_n$ with a proposition A that they justify. These formal structures represent what we will conceive of as different fragments. "$E \mid\!\sim_{\leq_i} A$" is true iff in all most $\leq_i$-plausible worlds in which E is true, A is true. $\leq_i$-plausible worlds are those worlds that are the most plausible according to the plausibility ordering $\leq_i$. If "$E \mid\!\sim_{\leq_i} A$" holds, then one can say that A is justified by E in the fragment corresponding to the ordering $\leq_i$.

Justification relations can be thought of as a qualitative way of modelling (Bayesian) absolute confirmation, as opposed to (Bayesian) incremental confirmation. By being fragment relative, different justification relations will underlie different plausibility orderings which will make it possible that a certain justification relation may hold between two propositions E and A in some fragment, but not in a different one.

Note a few things about this account.

This account of justification can be thought as a hybrid account of justification. Evidence plays an important role, so it has an evidentialist dimension, but evidence is not all that is required for a proposition to be justified. What makes a proposition justified is whether or not the evidence stands in an appropriate relation to that proposition[8].

Then, note that what is fragment-relative, here, are justification relations $\mid\!\sim_{\leq}$, and not propositions E or A. This is one aspect that significantly distinguishes this account from fragmented belief accounts, where beliefs

[7]Plausibility orders are pre-orders that are reflexive, transitive and total. For the present purpose, an intuitive understanding of these orderings will be sufficient. These kinds of orders are taken from two sources of inspiration. The first is the AGM belief revision literature (cf. Alchourrón, Gärdenfors, and Makinson, 1985, Grove, 1988), where it has been shown that the classical belief revision operator, '*', can be represented via total plausibility pre-orders. The second source is the literature on non-monotonic reasoning surrounding Kraus, Lehmann, and Magidor (1990) and Lehmann and Magidor (1992).

[8]The account is compatible with foundationalism, because it assumes that some propositions are more foundational than others; some serve the role of justifying another proposition, and others are in need of justification. It is also compatible with a version of coherentism, since it requires local (intra-fragmentary) coherence.

– modelled as propositions – are fragment relative, and where fragments are unordered subsets of W (sometimes also partitions of W). Here, instead, fragments are plausibility orderings over W, and propositions are not fragment-relative, but instead, justification relations are.

Next, because the justification relations are non-monotonic, it can happen that "$E \mid\!\sim_{\leq} A$" holds, while E, taken together with, say E', renders the justification relation invalid. This "new" piece of evidence E' would, then, be called a "defeater". The possibility of this happening is formally represented via a non-monotonicity property of the justification relation.

I assume non-monotonicity because I assume justification to be defeasible. One of the main drivers for accounts of fragmented belief, and the account of fragmented justification I develop here, is the desire to provide models of doxastic attitudes that remain attainable for real agents while remaining normatively significant. I take defeasibility to be an essential aspect of justification in science and "in real life", which is why I incorporate it into the account. Unlike defeaters, there can also be neutral pieces of evidence, E'' that could be added to a valid justification relation, $E \mid\!\sim_{\leq} A$, such that the relation remains valid, $E, E'' \mid\!\sim_{\leq} A$, but the neutral piece of evidence E'' by itself would not suffice to establish the relation, i.e., $E'' \mid\!\sim_{\leq} A$, would not hold. A neutral piece of evidence could be a proposition that is about something entirely unrelated, or a tautology. Whether or not adding a piece of evidence changes the validity of the justification relation depends on the specifics of the plausibility pre-order.

Moreover, intuitively, a fragment / plausibility ordering over W can be thought of as the background information that an agent is operating on in a specific circumstance. Given different background information (may include things like methodological standards, sources of information, stakes, research data on a specific subject matter that is salient/known to the agent etc.), I will argue that different propositions can become more or less plausible in different fragments.

Next, accounts of fragmented belief usually take fragmentation to have to do with lack of information. Note that here, the plausibility orderings always order all possible worlds, W (they are total). In this sense, infor-

mation is not restricted here, because each fragment orders not over proper subsets of W, but over the entirety of W. The reason for this modelling choice is that the preface paradox will be a case where fragmentation is warranted, but not because there is a lack of explicit information for the author. For the present account, information is restricted in a different way, by allowing for different background information to be operated on in different fragments. This means that it is not the set of propositions available to the agent that changes from fragment to fragment, but what changes is how plausible those propositions are, given different sets of (non-explicit) background information.

The last point is that, for this account to remain normatively significant there have to be some rationality constraints on fragments: Each fragment needs to remain consistent and closed under logic (i.e. adhere to traditional rationality norms), while cross-fragment inconsistencies are allowed and closure under logic does not have to hold across fragments. This means that it can be rational for some proposition A to be well-justified with respect to some fragment, but not with respect to some other fragment. This would not be allowed using traditional rationality norms.

Probably the most important question about this account is how the different plausibility orderings come to exist and how it can be rational to have multiple different plausibility orderings over the same set of worlds, W. In other words, why would it be rational to operate on multiple different sets of background information? I will answer this with the help of the preface scenario.

3.4 The solution

I will now present how the account of fragmented justification solves the paradox. I will first present the solution and then explain how it works.

The account of fragmented justification solves this paradox by relativising justification to fragments – a preface fragment, $\leq_P$, and a body fragment, $\leq_B$. Within the preface fragment $\leq_P$, there is a valid justification

relation in support of the preface proposition $\neg(A_1 \wedge ... \wedge A_n)$ that is based on the evidence E the author has in support of the preface proposition: $E \mid\!\sim_{\leq_P} \neg(A_1 \wedge ... \wedge A_n)$. The evidence for the preface proposition is that the author has published books in the past and each time, mistakes have been discovered later, or, that the author is fallible. Within the body fragment, $\leq_B$, there is a valid justification relation in support of the propositions $A_1, ..., A_n$: $E' \mid\!\sim_{\leq_B} A_1 \wedge ... \wedge A_n$. The evidence E' supporting those propositions is research the author has conducted on baby nutrition.

Note that unlike $A_1 \wedge ... \wedge A_n$ and $\neg(A_1 \wedge ... \wedge A_n)$, E and E' are consistent – the proposition expressing the fact that the author is fallible / has made mistakes in the past does not directly contradict any specific result of a scientific experiment on baby nutrition.

Note also, as mentioned, that within fragments, closure under logic holds. This is why, for the body of the book, the author has justification for each individual A_i – for each A_i there is a justification relation $E^i \mid\!\sim_{\leq_B} A_i$, and by conjunction introduction: $E' \mid\!\sim_{\leq_B} A_1 \wedge ... \wedge A_n$[9].

Because the justification relations are relativised to two different fragments, the preface proposition, $\neg(A_1 \wedge ... \wedge A_n)$, and the body propositions, $A_1, ..., A_n$, are not justified within one and the same fragment. For this to be the case, the preface proposition, $\neg(A_1 \wedge ... \wedge A_n)$, is not justified within the body fragment, while the body propositions, $A_1, ..., A_n$, are not justified within the preface fragment. Since closure properties do not apply across fragments, the inconsistent conjunction of propositions, $\neg(A_1 \wedge ... \wedge A_n) \wedge A_1 \wedge ... \wedge A_n$, is not justified within any one fragment, while the conjuncts, each are justified relative to a specific fragment. This solves the paradox.

But there are a few questions that have to be addressed to make this a convincing solution:

1. Why is it rational to entertain a preface and a body fragment, i.e. two different plausibility orderings over W?

[9]Strictly speaking, some monotonicity (which does not hold generally) is needed here, too. We can put together the pieces of evidence $E^1, ..., E^i, ..., E_n$ together by conjunction to result in E'.

2. Why is it that the preface and the body propositions cannot be justified within one single fragment?

3. Why can't one get rid of the fragments to result in an overall consistent doxastic state?

4. There should be limits to fragmentation on pain of irrationality. What should those limits be?

5. Is this an account of rationalising mistakes?

6. How does this solution differ from fragmented belief solutions?

I will now go through each of these questions and answer them.

Question 1 Why is it rational to have a preface and a body fragment, i.e. two different plausibility orderings over W? I will now argue for why it is rational to introduce a preface and a body fragment in a situation like this.

First of all, before arguing why this specific way of fragmenting information is rational, let me argue for why fragmentation makes sense in this kind of scenario in general. Fragmentation can be viewed as a strategy to deal with overall inconsistent information where there is no prima facie reason to give up either piece of information. We should not give up the preface proposition, because it is well justified, but we also should not give up the body propositions, because each of them is well justified. Given such a situation, traditional methods for dealing with inconsistent information, such as suspension of judgment or belief revision, are unsatisfactory, because they would force one to give up some justified proposition to rectify the situation. But there is no reason to give up either proposition.

Fragmented justification is a new way of dealing with such situations that reflects that the reason the situation arose is that each proposition is well justified, and that allows to keep all propositions within one's doxastic state, while getting rid of the outright inconsistency – by relativising justification to fragments.

Now, why does it make sense to split into a preface and a body fragment with respect to justification? There are a few different factors that make it rational to fragment in this situation, i.e. that make it rational to have two different plausibility orderings over W. These are the factors that guide the fragmentation process: Subject matter of the pieces of evidence, sources of information, the stakes, the type of questions it is appropriate to ask about any given proposition, as well as uncertainty / defeasibility. These factors and how their manifestations differ for the preface and the body parts of the book will justify operating against different sets of background information in each part of the book. And they will be the reason why it is rational for the author to have two different plausibility orderings over the same set of possible worlds W, and thus, why it is rational to fragment.

These are guiding principles for the individuation of fragments which I will go through now in the context of the preface paradox as a case study. The goal for further work will be to develop an explicit set of norms for the individuation process of fragments, i.e. when and how an agent should fragment, generally and independently of the preface scenario.

First, for the body part of the book: The subject matter of the evidence E' is baby nutrition – specifically results of scientific experiments and the like that are able to support the body propositions $A_1, ..., A_n$. The sources of the information in the body of the book are scientific data, such as experiments. Furthermore, for the body of the book, it is appropriate (for the author or reader) to ask questions about the support / justification of the claims the author makes in the book. This type of justification can be interpreted as first-order justification – it is directly about the subject matter of the book – baby nutrition. Furthermore, in general in data based sciences, but especially here, because the number of propositions $A_1, ..., A_n$ communicated in the body of the book is large, there is uncertainty involved. Each proposition is well-justified, but justification is defeasible. It might be that in some time, there author will learn there to be a defeater to undermine the justification of one of the propositions $A_1, ..., A_n$, although all of them are perfectly justified right now.

Let us turn to the preface part of the book: The subject matter of the

evidence E justifying the preface proposition $\neg(A_1 \wedge ... \wedge A_n)$ is not baby nutrition as in the body of the book, but the author's track record with regards to making mistakes / her fallibility. It is not any specific result of a scientific experiment or the like, as it was in the body of the book. The source of information, here, is not any particular piece of scientific evidence that may defeat a justification relation the author establishes in the main part of the book, it is rather, inductive evidence in terms of her track-record with regards to previously published books / her fallibility. Here, the stakes change for the author in comparison to the body part of the book. She takes a step back and becomes more cautious about the claims she makes in the body part of the book, because she is aware of the defeasibility of justification, which is why she communicates the preface statement.

Furthermore, for the preface part, it is not appropriate (and does not make any sense) to ask about the support / justification of any individual claim made in the book – such as "Bananas are good baby food". The preface fragment is not the place where the reader learns about baby nutrition and the (first-order) justifications the author provides in support of her claims – the reader learns about this in the body part of the book. In relation to the preface, the reader may ask about the support / justification for the preface proposition. This justification can be interpreted as a kind of second-order justification, because it is not any individual piece of scientific evidence about baby nutrition as is the first-order justification in the body of the book. Rather, the piece of evidence justifying the preface proposition is about the author's track record / fallibility more generally. The justification relation that will validate the preface proposition within in the preface part of the book cannot undermine or defeat any particular justification relation established in the body of the book, since it is justification of a different type. And again, uncertainty is relevant in the preface as well. Because the author is aware that justification is defeasible, and because she is aware that she is fallible, she comes to be justified in the preface proposition.

One might think – why does it matter that in the preface and the body,

respectively, there are different subject matters, different sources of information, different stakes, different types of evidence relevant, and different questions that are appropriate to ask. The answer to this question is that, because this is the case, it becomes rational that the author communicates an overall inconsistent set of propositions. It is because of these differences that such a situation can be rational. These differences make it rational to operate upon different sets of background information in the two different fragments. These two sets are not compatible with each other, but both are equally legitimate. Neither set of background information is better or worse. But they are different, different enough to warrant fragmentation. If all propositions involved had the same subject matter, the same source of information, the same types of evidence involved, the same types of appropriate questions, and the stakes were equal, then why should one fragment? In such a case, it would be very difficult to rationalise fragmentation, and overall inconsistency.

These differences that are present here are factors that make it possible to grant the rationality of inconsistencies, because they can be sources of rational incoherence or inconsistency. But one might ask: Why do these factors make fragmentation and overall inconsistency rational? The reason is uncertainty. Uncertainty makes this situation rational. Uncertainty is the reason why, given two different sources of information, different subject matters, different stakes, fragmentation is warranted here.

What should one do in such a case? Randomly give up a piece of perfectly fine information (as orthodox epistemology would recommen), or fragment until one knows more? One should fragment to be able to keep the information that is well-justified. Forcing consistency without fragmentation would make one lose perfectly fine information which would be epistemically defective. Fragmentation provides an opportunity to figure out why the inconsistency arose in the first place, keep it if it is warranted or eventually get rid of it if a defeater appears. But as long as there is no defeater currently available, the author should keep all her information and go fragmented. Of course, if there were only one source of information, this source should better provide consistent information, otherwise one would

doubt its trustworthiness.

Likewise, as it is the case here, it can be that an inconsistency arises because the types of justifications and underlying subject matters (and order type, first vs. second order) vary. Should one give up randomly a perfectly fine piece of information or should one fragment? One should fragment. Of course, if there were no different subject matters and different orders (first vs. second order) of justification involved, then it would be harder to argue why one should go fragmented.

In the body and the preface, the author operates on different background knowledge; she uses different sources of information, the propositions relevant to her have different subject matters and different orders (first vs. second order), and she has different stakes associated with the preface and the body part of the book, respectively. Testimony of these differences is that there are different questions that one may appropriately ask in the preface and the body of the book, respectively.

The different background knowledge the author operates on is reflected in two different plausibility orderings / fragments. In the body fragment, the claims the author makes in the body of the book and their respective justifications have high plausibility, or, it is highly implausible that any one of the claims she makes in the body of the book is false (i.e. that the preface proposition is true). Why is this the case? Because, as Makinson (1965) states in the original formulation, each claim the author makes in the body of the book is highly justified. On the basis of the valid justification relations in the body of the book, the preface proposition is not very plausible. Formally, this means that in the body of the book, all worlds in which all $A_1, ..., A_n$ are true are more plausible than any world in which (at least) one proposition of the $A_1, ..., A_n$ is false.

In the preface fragment, taking on a different perspective; taking a step back, because of the author's track record / fallibility, it gets highly plausible that there is some claim in the body of the book that is false. But the justification for this claim is not any particular piece of evidence that directly defeats a claim she makes in the body of the book, but her track record and / or general human fallibility – i.e. second-order justification. It

is very important here, that the evidence justifying the preface proposition cannot undermine any of the valid justification relations established in the body part of the book.

Formally, this means that in the preface, all worlds in which (at least) one of $A_1, ..., A_n$ is false are more plausible than any world in which all $A_1, ..., A_n$ are true.

Question 2 Another point that has to be made clear for this solution to be convincing is why it is not the case that $\neg(A_1 \wedge ... \wedge A_n) \wedge A_1 \wedge ... \wedge A_n$ can be justified within one and the same fragment. Because if this were the case, the solution would not work.

Recall first that pieces of evidence are not fragment-relative on this account. Within the preface fragment, what actually holds is $E, (E') \ \mid\!\sim_{\leq_P} \neg(A_1 \wedge ... \wedge A_n)$. E is about the author's track record / general human fallibility, while E' comprises the scientific evidence the author has for her claims. The scientific evidence on baby nutrition, E', does not suddenly become unavailable or implausible within the context of the preface. But E' plays the role of a neutral piece of evidence in this relation. It does not invalidate the justification relation for the preface paradox, but it cannot justify the preface proposition by itself: $E' \not\mid\!\sim_{\leq_P} \neg(A_1 \wedge ... \wedge A_n)$. E' – the justifications for the claims about baby nutrition – does not undermine the preface proposition directly, however, it cannot support it on its own, because it is not of the right type to be able to support the preface claim. Scientific evidence on baby nutrition does not justify the preface proposition. But E' is still accessible (and plausible) in the preface fragment. This shows that the availability of evidence to the agent is not bound to a particular fragment, but rather, that, what changes is the "impact" evidence can have on justification relations, because their plausibility rank may change in different fragments.

Similarly, in the body, what actually holds is: $(E), E' \ \mid\!\sim_{\leq_B} A_1 \wedge ... \wedge A_n$. But again, $E \ \mid\!\sim_{\leq_B} A_1 \wedge ... \wedge A_n$ alone does not hold. The reasoning is similar: The fallibility / track record of the author, E, by itself cannot support the claims the author makes in the body of the book, because it is not of the

right type. What is needed to support the claims in the body are, for example, the results of scientific experiments. But E is still accessible (and plausible) in the body fragment – the availability of evidence to the agent is not bound to a particular fragment. Pieces of evidence are available in any fragment – it is just that different pieces of evidence can support different claims in different fragments.

Based on this, one might object that if E and E' are in principle accessible in all fragments, then it would also hold that $E \mathrel{|\!\sim_{\leq_B}} \neg(A_1 \wedge ... \wedge A_n)$. This would be a problem, since then, by closure under conjunction within the body fragment, the paradox could not be solved since, then, what would also hold is: $E, E' \mathrel{|\!\sim_{\leq_B}} A_1 \wedge ... \wedge A_n \wedge \neg(A_1 \wedge ... \wedge A_n)$ (likewise for the preface fragment), and this would mean that the account of fragmented justification cannot solve the paradox.

But $E \mathrel{|\!\sim_{\leq_B}} \neg(A_1 \wedge ... \wedge A_n)$ does not hold. The reason for this is that according to the plausibility ordering of the body of the book, any world in which the agent is fallible, i.e. where E is true, and where all $A_1, ..., A_n$ are true is more plausible than any world where E is true and the preface proposition $\neg(A_1 \wedge ... \wedge A_n)$ is true. Because the truth conditions of $\mathrel{|\!\sim_{\leq}}$ depend on truth in all most plausible worlds, $E \mathrel{|\!\sim_{\leq_B}} \neg(A_1 \wedge ... \wedge A_n)$ does not hold, since worlds in which E and $\neg(A_1 \wedge ... \wedge A_n)$ are true are not among the most plausible worlds. Based on the strong justification in the body of the book, it is more plausible that the author is fallible and did not make a mistake than that they are fallible and did make a mistake. It is only when one switches to the preface fragment, and thereby changes the perspective from individual justification relations about baby nutrition to general human fallibility, and thereby employs a different set of background information, that worlds in which the author is fallible and one of $A_1, ..., A_n$ is, in fact, false, are more plausible than worlds in which all $A_1, .., A_n$ are true.

The justificatory "labour" is done in the preface by E (second-order justification – the author's track record / fallibility) and in the body by E' (first-order justification – results of experiments, scientific data). The author's fallibility can neither justify nor invalidate any of the individual

claims about baby nutrition she makes in the book, nor can the evidence for the claims she makes in the book justify or invalidate the preface proposition.

Question 3 Why can't one get rid of the fragments to result in an overall consistent doxastic state? One might argue that it must be possible to dissolve the fragment-relativity of justification relations in some way, which could be a threat to this solution to the paradox.

It is known from social choice theory, in particular, from Arrow (1950)'s impossibility theorem, that (rational) aggregation of preferences is not possible given a number of prima-facie plausible rationality postulates[10]. In the judgment aggregation literature (cf. Dietrich and List, 2016) the same has been shown for the aggregation of propositions. This makes it plausible to suspect that an instance of the same general result might hold for the aggregation of fragments and their underlying plausibility orderings (total pre-orders) in the context of fragmented justification. Generally, if one would like to merge different fragments / plausibility orderings nonetheless, one will have to give up on certain prima-facie plausible rationality criteria.

But one might object that in the specific situation of the preface paradox, it could be possible and useful to find a suitable merging mechanism to get rid of fragmentation. Note that this would not necessarily be an objection to a fragmented justification solution to the paradox – just because fragmentation might be dissolvable, this does not render fragmentation per se wrong, nor a fragmented justification solution to the preface paradox uninteresting.

I will now go through some options for merging the two fragments $\leq_P$ and $\leq_B$ in the preface scenario. Recall that the plausibility ordering for the preface fragment ranks worlds in which the preface proposition is true the most plausible and worlds where all $A_1, ..., A_n$ are true, and the preface proposition false, as the least plausible. The plausibility ordering for the body fragment ranks the worlds where all $A_1, ..., A_n$ are true the highest

[10]Such as anonymity, independence or universal domain. For an introduction to judgment aggregation, cf. List and Polak (2010).

and worlds in which the preface proposition is true the lowest. In order words, the respective plausibility orderings for the preface and the body of the book are mirror-opposites in terms of the ranking of worlds where the propositions $A_1, ..., A_n$ and the preface proposition $\neg(A_1 \wedge ... \wedge A_n)$ are true, respectively[11].

The first potential option for merging the two fragments would be to take the union of the two plausibility orderings. The first problem concerning this option is that it is not clear what the union(-order) of two pre-orders is supposed to be. Nevertheless attempting to take the union-(order) of these orderings would result in there not being any gradation of plausibility at all anymore – and thus there not being any interesting plausibility ordering anymore, which would be a very uninformative outcome of the aggregation process[12].

Alternatively, one might want to merge by taking the intersection-(order) of both orderings. Unlike for the union, the intersection of two pre-orders is a defined notion, but the resulting order will only be partial. However, nevertheless this option is in principle possible here, the result would be that each world would be equally plausible with itself, but that there would be no overall ranking across the different worlds anymore. Again, this seems like an undesirable result of an aggregation process[13].

Majority and unanimity aggregation procedures seem unfitting as well, because what makes fragmentation attractive is that the plausibility orderings may look very different for such rules to be very difficult to apply. And specifically, for the preface, since there are only two fragments, both would need to look at least very similar for a majority or unanimity aggregation

[11]It could also be that the orderings are not completely opposed which would not significantly change how this account solves the paradox, but would open up more options for the aggregation of fragments. For expository purposes, I show how the aggregation options turn out with the ordering explained above.

[12]For the interested reader: Let $\leq_1$ be $t \leq u \leq v \leq w$,and let $\leq_2$ be $w \leq v \leq u \leq t$. Then, the union would result in $t = u = v = w$.

[13]For the interested reader: Let $\leq_1$ be $t \leq u \leq v \leq w$,and let $\leq_2$ be $w \leq v \leq u \leq t$. Then, the intersection would result in each A, B, C and D being equally plausible with itself, respectively, but there not being any overall plausibility ranking between t, u, v and w.

procedure to be applicable, which is not the case here.

Another option could be to use a variant of the Borda count (cf. eg. Pacuit, 2019), by trying to result in an overall unfragmented doxastic state through computing mean plausibility values for each world, across fragments. However, this option does not work for pre-orders, since there are no "plausibility values" attached to the different worlds in the orderings. One might think that, one could still introduce values through coding orders through numbers that one attaches to the worlds in the ordering. But again, those numbers would be merely ordinal, and taking an average or mean is not defined for an ordinal sequence like a pre-order.

After this brief survey, existing aggregation methods seem unpromising for the preface paradox scenario. But, even if there were a suitable way to merge fragments, it will not always be necessary. If the fragmented belief state were to pose no problems for epistemic or practical rationality, why should one get rid of fragmentation, if it is epistemically warranted?

In the preface case, it seems that the author's fragmentation would not necessarily pose any problems for epistemic or practical rationality. The author is epistemically rational in entertaining the set of overall inconsistent propositions within a fragmented doxastic state, and there seems to be no problems for practical rationality either. Imagine a potential reader of the book on baby nutrition – the fact that the total set of propositions contained in the book is inconsistent is not a problem for her understanding of the contents of the book. Through reading the book, she learns about the author's contributions in the field of baby nutrition while having in the back of her mind that the author is fallible and that one of her claims might be proven wrong later down the line. This is a perfectly fine situation for the reader to be in, both in terms of epistemic as well as practical rationality.

Thus in the context of the preface paradox, one cannot sensibly and it seems not necessary to de-fragment. Generally, and independently of the preface paradox, the topic of (and the need for) de-fragmentation of justification fragments is an interesting topic that should be further investigated.

Question 4 There should be limits to fragmentation on pain of irrationality. What should those limits be?

One might think that, maybe, the author should fragment more-fine-grained than merely into a preface and a body fragment. One might want to introduce many more fragments in the body of the book, corresponding to chapters, sections or even paragraphs. One might argue for this by saying that chapter 1 is, say, about baby nutrition for newborns, chapter 2 about baby nutrition for 3-12 month-old's and so on. While the overall subject matter of the body of the book would still be baby nutrition, this could be made more fine-grained by distinguishing into different sub-subject matters.

Let me explain why this would not be a good idea. Firstly, note that dividing the body up into not one but many fragments would not change anything about the preface fragment being an entirely distinct fragment from any body-fragment(s).

Secondly, and more importantly, introducing more fragments than the two already introduced would not be a good idea. The reason for this is that fragmentation comes with benefits and costs and the more fragments are introduced, the more the costs will outweigh the benefits. One benefit of fragmentation of justification is that it is able explain why the author is rational in entertaining an overall inconsistent set of propositions. It allows agents to retain and sort overall inconsistent information within their doxastic state without immediately being required to get rid of it.

However, fragmentation does come with costs. First, the author as well as a potential reader of the book would most likely want to be able to draw inferences within the book. Since inferences can only be drawn within but not across fragments, the more fragmented the author becomes, the more difficult it will be to draw inferences and conclusions from the book. Moreover, each fragment will correspond to a different plausibility ordering and that comes with its own fragment-relative justification relation. If there were very many fragments, one might worry that both the author and any potential reader would not be cognitively able to track all of these different plausibility orderings and justification relations. And generally, the more

fragmented the author becomes, the more and more difficult it becomes for them to keep track of the fragmented information. If the author is very or even extremely fragmented, such costs will outweigh the benefits of fragmentation. Because of these costs, the author should introduce as few fragments as possible. In the case of the preface paradox, keeping the number of fragments at two is the sweet spot. A more general question worth investigating would be whether it is always possible to determine an ideal number of fragments.

This discussion on limits of fragmentation ties in with the question whether there should inter-fragmentary rationality criteria for fragments. In other words: "How should fragments relate to each other?" It is an essential aspect of the account that there may be rational inconsistencies across fragments. However, one might still want to have some limits on how different the fragments may look – especially if there are many. One option for such a limit could be to stipulate a maximal distance measure between plausibility orderings. For an account of fragmented justification to be applicable to a wide variety of situations, while remaining normatively significant, it will be important to decide how exactly to choose such a measure. Such a measure will be part of a general set of norms for the individuation process of fragments to be developed in further work.

Question 5 One might think that this account is an account of rationalising mistakes. But that would be bad, the account should definitely not rationalise any kind of situation. If this were the case, then the account would miss its goal of being normatively significant, and of capturing something essential about the concept of rationality.

An obvious mistake that this account should not rationalise would be this: The author publishes a book containing only one sentence expressing one proposition A_1 in the body of the book and in the preface, she communicates the sentence expressing the corresponding preface proposition $\neg A_1$. The account should not rationalise this single proposition preface case, since a book containing only A_1 and $\neg A_1$ seems very problematic for both epistemic and practical reasons.

But, one might think that she has good justification in the hypothetical situation just presented: She has good justification for the claim A_1 expressing, eg., that bananas are a good baby food, and that she has good justification for the preface proposition, i.e. past mistakes in books or general human fallibility. And, similar to the original preface case, these justifications remain of different types – first vs. second-order.

However, there are two important differences between the original paradox and the single-proposition case. First, in the original preface paradox, the preface proposition $\neg(A_1 \wedge ... \wedge A_n)$ does not directly negate any particular one of the individual claims $A_1, ..., A_n$ she makes in the body of the book. It only negates the conjunction of these claims. However, in the single-proposition scenario, the new preface proposition $\neg A_1$ does directly negate the claim the author makes in the body of the book. This makes the single-proposition case much stronger and it becomes harder to argue for its rationality.

The second difference is that, actually, in the single-proposition case, the preface proposition is no longer justified. The justification of the preface proposition crucially relies on the fact that there are many propositions $A_1, ..., A_n$ postulated in the body of the book (which also increases the amount of uncertainty the author faces in that case). If the author only postulates one single proposition, human fallibility and / or having made mistakes in the past would not be strong enough to justify the claim that the proposition A_1 for which the author has excellent first-order evidence is false. It only becomes strong enough when the author is postulating many propositions in the body of the book, since this makes the author more fallible (because there is more uncertainty involved) than in the single proposition case. So, what counts as "strong enough" varies in different circumstances, and this decision is also part of the background information that determines the plausibility orders, and thereby, which propositions are amongst the most plausible worlds, and thus, "strong enough" to justify another proposition.

Although justification is defeasible, and humans are always fallible, an account of justification that renders the preface proposition justified in the

single-proposition case would be too weak. Because the justification for the preface proposition is lost in the single-proposition case, the preface proposition is not justified, which is why the paradox would not arise, since for it to arise both the preface and the body proposition would need to be well justified, which is not the case here.

Question 6 The last question that needs to be answered is how this does solution differs from fragmented belief solutions.

This account of fragmented justification is closely related to accounts of fragmented belief in that both are fragmentation accounts. However, there are significant differences. For the account of fragmented justification, justification relations are fragmented and fragments are total preorders on possible worlds. For accounts of fragmented belief, beliefs are usually modelled via propositions. Those propositions are then taken to be fragment-relative. For fragmentation accounts of belief, fragments are proper (unordered) subsets of the total set of possible worlds W (sometimes also partitions of W). So, both the definition of a fragment differs for both accounts, as well as what "things" are taken to be fragment-relative.

The account of fragmented justification I develop here does not take propositions to be fragment-relative. Propositions are available in all fragments. What differs from fragment to fragment is how plausible they are. But what differs from fragment to fragment for fragmented belief accounts is which propositions are available. I will shortly argue that fragmented justification provides a rationality argument for fragmented belief, and that, therefore, fragmented belief accounts should take on board a fragmented reading of justification. This would combine the way in which belief is traditionally taken to be fragmented with this account of fragmented justification, and result in a mix between the two. But given accounts of fragmented belief how they appear in the literature, they are very different from the account of fragmented justification developed here.

And lastly, note that this account is silent on whether belief should be fragmented or not, it does not talk about belief at all, but rather about whether or not a proposition is well justified or not. I will come back to

this at the very end of the chapter.

By addressing the above questions, I have argued for why the solution that the account of fragmented justification provides is convincing. It is rational for the author to fragment her information in a way that gets rid of the inconsistent conjunction of propositions. This solution of fragmented justification can explain the appeal of the paradox. The author is rational, because she has good justification for the body propositions as well as the preface proposition. However, since the preface proposition and the body propositions are not supported by a justification relation within the same fragment, no inconsistency arises. The fragment-relativity of the justification relation blocks the inconsistent conjunction of propositions being justified within one and the same fragment. She does not have to arbitrarily give up any proposition, but fragmented justification allows her to keep them all within her doxastic state while getting rid of the problematic inconsistency.

3.5 Evaluation

In this short evaluation section, I focus on the two big advantages that this solution has to other solutions to the preface paradox, especially to fragmented belief solutions.

3.5.1 Justification at the core

The first deficiency of existing solutions to the preface paradox is that none of them focus on justification as the central concept responsible for the paradox. As Makinson (1965) stresses, it is because of the justification that each proposition involved in the paradox enjoys, that it seems that the author is rational. Solutions that specifically use accounts of fragmented belief also do not reflect the importance of the justification of the propositions involved in the paradox.

The account of fragmented justification gets rid of this deficiency that

other existing solutions face by providing a solution that uses the central concept of the paradox – justification. It shows how it can be the case that both $\neg(A_1 \wedge ... \wedge A_n)$ as well as $A_1, ..., A_n$ are all well-justified while getting rid of the outright inconsistency by relativising justification to fragments. I have provided reasons for why it makes sense to relativise justification to fragments, and I have showed how this provides a novel solution to the paradox.

3.5.2 Boosting credibility of fragmented belief

Secondly, fragmented belief solutions need to answer the question whether and why fragmented belief can be rational. The second perk of the account of fragmented justification is that it provides a rationality argument for fragmented belief accounts and therefore increases the attractiveness of fragmented belief solutions to the paradox. The reason why fragmented belief solutions make sense, and more generally why fragmented belief can be rational, is because justification is fragmented, too.

Let me explain why this is the case. Recall the reasoning for why fragmented belief accounts argue that belief is fragmented in the preface case: The preface and the body of the book belong to two different fragments because the propositions in the preface and the body of the book are about different subject matters. But how does the author get to believing these propositions in the body and the preface in the first place? She does so on the basis of her justifications, as Makinson (1965) states in the original formulation of the paradox. One reason why fragmentation can rationally arise on the level of belief is if this fragment-relativity is already present on the level of justification. What renders a belief rational is its justification. Each belief in a belief state is rationalised by its respective justification. If this belief state is fragmented, what rationalises this fragmentation is the fragmentation of its underlying justifications.

This suggests the following conclusions with respect to the more general relationship between belief and justification. If one wants to be fragmented with respect to belief, one might want to be fragmented with respect to jus-

tification, too[14]. Because the account of fragmented justification provides a rationality argument for fragmented belief accounts, it manages to boost the credibility of fragmented belief solutions to the paradox, in addition providing a novel and attractive solution on its own.

Finally, also in light of the relation between fragmented justification and fragmented belief just touched upon, fragmented justification is compatible with two possible stories to tell about what the author communicates with the reader, because it is silent on whether the justified propositions $A_1, ..., A_n$ and $\neg(A_1 \wedge ... \wedge A_n)$ are (fragmented) beliefs or not.

The first story is to accept fragmented belief in addition to fragmented justification, and conclude that the author communicates fragment-relative justified beliefs, $A_1, ..., A_n$, about baby nutrition as well as, in a separate fragment, about her fallibility, $\neg(A_1 \wedge ... \wedge A_n)$. In this case, the fragmented justification solution provides the additional insight to fragmented belief solutions that it explicitly models the justifications involved which are the reason for why the paradox arises. It also explains why the fragmentation of belief arises in the first place, and moreover, why the author is rational in entertaining this fragmented belief set.

The second story is that the author actually does not communicate beliefs at all in her book. What she communicates, rather, are that her claims $A_1, ..., A_n$ and $\neg(A_1 \wedge ... \wedge A_n)$ are well supported / justified relative to the respective fragments. This second story leaves open whether belief simpliciter should be fragmented or not[15].

[14] An interesting question to investigate would be whether the other direction holds, i.e. whether one could adopt a fragmentation account of justification without adopting a fragmentation account of belief.

[15] One might worry that that if the author claims but not believes the propositions $A_1, ..., A_n$ she expresses in the book, then she would be entirely insincere. But, the notion of "claiming" that would support this worry would be very strong, since her claims are very well justified. Therefore, this does not render the author insincere. Whether she publishes a claim or not depends on whether it is justified or not. With the publishing of the well-justified propositions, one might interpret this as the author communicating her acceptance, instead of her beliefs. For accounts of acceptance cf. eg. Van Fraassen (1980), Stalnaker (1984), L. Cohen (1989), L. Cohen (1992).

3.6 Conclusion

In this chapter, I sketched an account of fragmented justification according to which justifications are valid fragment-relative justification relations between propositions. I showed how this account solves the preface paradox. Fragmented justification is an attractive way of dealing with an overall inconsistent set of propositions, where, crucially, each proposition is well-justified, where one does not have to arbitrarily give up some proposition. Instead, by relativising justification to fragments, the author can retain all her well-justified information.

Because closure under conjunction does not apply across fragments, the inconsistent propositions are only justified relative to their respective fragments, but not within one and the same fragment. Thereby, the author remains overall rational in publishing a book containing both $\neg(A_1 \wedge ... \wedge A_n)$ and $A_1, ..., A_n$.

Furthermore, this solution shows why it can happen that inconsistent propositions are each well-justified: The reasons are difference in sources of information, subject matter, stakes, types of justification and types of questions appropriate to ask in the preface and the body of the book, respectively, as well as uncertainty.

This solution is novel, because the account of fragmented justification employed here is novel. I provided two more reasons for why this solution is attractive. First, unlike existing solutions, it employs the central concept in the paradox – justification. And second, fragmented justification can increase the credibility of fragmented belief solutions.

One may view this first chapter as a case study that is the starting point for a fuller development of an account of fragmented justification. As explained, the preface case is a case where it is rational to fragment with respect to justification and where de-fragmentation would not make much sense. Part of a full account of fragmented justification will be the development of a more general set of rationality norms that guide both the individuation of fragments (When and how should one fragment?) as well as the merging of fragments (When and how should one get rid of fragmen-

tation?). Specifically, it will be interesting to investigate how fragmented justification applies to different cases of overall inconsistent sets of propositions – not every inconsistency should be rationalisable by this account of fragmented justification.

Subject to further work in the following chapters will also be the development of a Bayesian account of fragmented justification, as well as to draw further conclusions regarding the relationship between belief and justification with respect to fragmentation more generally.

Chapter 4

Fragmented Justification

This chapter presents a (fuller development of a) fragmented account of justification. The account is motivated via an example from climate science where, within the IPCC assessment report for climate change, an overall inconsistent set of justified beliefs seems to be asserted. The main reason why this happens will be the extreme uncertainty involved in learning about complex systems such as the climate. The chapter explains how the toolkit of orthodox epistemology is limited in and not suited for dealing with situations like this, and argues that fragmentation is a more appropriate strategy than existing approaches to deal with (seemingly rational) inconsistent doxastic states. After presenting the example, the chapter gives a formalisation of this example. The upshot of this formalisation is that justifications are represented via non-monotonic fragment-indexed justification relations that hold between propositions that need justification, and evidence. The chapter explains why fragmentation is rationally warranted here, and develops an explicit set of precautionary norms that should guide the fragmentation process.

It, furthermore, argues that the account of fragmented justification developed here is compatible with fragmented belief, and how this account of fragmented justification can serve as a novel argument for the rationality of fragmented belief accounts developed in the literature.

And lastly, the chapter explains how sometimes, even though fragmen-

tation might be epistemically rational, de-fragmentation will become necessary in light of practical rationality. It shows how de-fragmentation can work for epistemically rational fragmented agents, which is a novel discussion within the literature on fragmentation.

Please recall, as already noted in the beginning of the dissertation, that a version of this chapter has, as of August 2025, been published with Synthese (Pillin, 2025).

4.1 Introduction

Since Davidson (1982b), Lewis (1982), Stalnaker (1984) and Cherniak (1986), fragmentation accounts of belief have become popular in epistemology. According to them, belief is fragmented in the sense that different beliefs an agent might have "live" in different mental fragments, as opposed to the more traditional view according to which all beliefs reside in one big "bubble". In each fragment, some normative closure conditions such as coherence or consistency and closure under logic apply, however they do not apply across fragments. In contrast, on the more traditional view, *all* beliefs of an agent ought to cohere or be consistent with each other.

This chapter develops an account within the spirit of fragmentation accounts of belief. However, the account is not of fragmented belief, but rather of fragmented justification. Fragmented justification is novel within the literature on fragmentation. Traditional accounts of justification, such as reliabilism or evidentialism, are all implicitly based on an unfragmented picture of rationality[1]. Thereby, this account expands fragmented accounts of doxastic attitudes, like belief, by a fragmented account of justification. Justification will be fragmented in the sense that the justification relation that holds between evidence and propositions that are being supported by that evidence is indexed to fragments. This will make it possible that some proposition might be justified in some mental fragment but not in another one.

[1] I provided an argument for this in the end of chapter 2.

The account of fragmented justification will be motivated and presented in the context of an example from climate science, where, within the IPCC assessment report for climate change, an overall inconsistent set of beliefs is asserted. This example poses a dilemma for orthodox epistemology, because orthodox epistemology would deem this as a case of irrationality, whereas the inconsistent belief set, intuitively, seems rational, because each of the beliefs involve is well justified, as I will show. Orthodox epistemology is not equipped with a strategy appropriate for dealing with the specifics of the epistemic situation at hand – which is mainly due to the severe uncertainty associated with learning about a complex system such as the climate – and it cannot explain the intuition that this is not a case of irrationality that needs to be somehow resolved or gotten rid of.

Fragmented justification, unlike orthodox epistemology, will be able to provide an appropriate solution to this dilemma, making use and providing a precise account of how the structure of the underlying justifications below the scientists' beliefs is fragmented, and rationally so.

This account is intended to be a normative guide for the rationality of agents, it is not to be understood as an account of rationalising "mistakes". The chapter provides reasons for why fragmentation can be rational, and furthermore, presents a set of precautionary norms that should guide the fragmentation process.

Moreover, I will discuss how fragmented justification relates to belief and argue that the account of fragmented justification developed here is compatible with both, a fragmented, as well as an unfragmented notion of belief. This is significant because it thereby allows for an agent's epistemic state to be fragmented at a lower level (justification), while being unfragmented on a higher level (belief). This discussion involves, furthermore, a novel argument for the rationality of fragmented accounts of belief that goes beyond and improves on the literature on the rationality of fragmented belief.

Lastly, I will also discuss how, for the sake of practical rationality (and unfragmented belief), de-fragmentation can become necessary, which is a novel discussion within the fragmentation literature. I show how there is

no generally applicable rule as to which aggregation procedure one should employ in any specific situation, and more generally, that in some cases, there will not be any sensible aggregation mechanism available. I argue that the aggregation of fragments is a situation of inductive risk, building on and providing a new application of the seminal inductive risk argument by Rudner (1953), and I will argue that the process of finding a way to aggregate fragments should be guided by non-epistemic value considerations about the risk of making a wrong call.

The chapter is structured as follows: Section 2 presents a motivating example for the account of fragmented justification from the field of climate science. In section 3, orthodox epistemology tries to solve the problem this example faces, but fails, and fragmented justification comes to the rescue. In section 4, I formally represent this example and explain what it means for a proposition to be justified on my account. Section 5 summarises why fragmentation is rational in this example. Section 6 presents a set of precautionary norms surrounding when and how one should and should not go fragmented. Lastly, section 7 is concerned with the relationship between justification and belief in the context of fragmentation, with special focus on how fragmented justification can serve as a novel and improved rationality argument for fragmented belief accounts in the literature, as well as with the topic of de-fragmentation in light of practical rationality.

4.2 A motivating example

This example is from the field of climate science. Climate science deals with understanding the complex system that is the climate, and representing that system (or parts of it) in a computational model to generate scientific hypotheses and predictions about the climate and climate change (cf. Parker, 2014).

Although there is much debate about what exactly constitutes a complex system, there is no debate that the climate is a complex system. Features of complex systems include that they consist of many interacting compo-

nents with a multi-layered structure. The climate consists of five major subsystems; the atmosphere, the hydrosphere, the cryosphere, the land surface and the biosphere, where each of these subsystems consists of sub-subsystems. These systems interact a lot; for example, the atmosphere and the ocean exchange water vapour and heat which leads to condensation, precipitation and supplies the weather systems with energy. They have a complex aggregate behaviour which is often non linear (and thus cannot be described via a set of linear equations). Examples of this are movements of large-scale weather systems.

The workings of the climate do not (only) depend on simple causal laws, but on local and highly fragile causal connections. This means that small local perturbations of the system can and do have large and unforeseeable consequences. Furthermore, the climate system is highly influenced by feedback processes that constantly change and make the system evolve (cf. Ladyman, Lambert, and Wiesner, 2013, IPCC, 2023).

Scientists model this complex system via mathematical equations with the help of computers and try to generate and estimate solutions and predictions about, among other things, future climate change.

There are two types of uncertainty influencing the representation process of the climate, and thus the predictive accuracy of the results of these representations: Uncertainty within the model and uncertainty about the model and its uncertainty itself. Parker (2014) distinguishes different varieties of the first kind: Structural, parametric and initial condition uncertainty. Structural uncertainty is about what form our modelling question it should take, for example, whether we should investigate precipitation in Europe or global mean temperature development. Parametric uncertainty is about the values that parameters should take in a model of the climate, while initial condition uncertainty is about which values one should assign to model variables at the start of the simulation.

The second kind of uncertainty is meta-uncertainty. This is, first, uncertainty about whether the model is informative about the climate at all, and secondly, uncertainty about one's object-level uncertainty estimates themselves.

These different kinds of uncertainty stem from the high complexity of the climate system: The scientist, albeit being an expert about climate, only has a very limited understanding of the actual processes constituting the climate. And a consequence of this is that each decision climate scientists make when modelling the climate and generating hypothesis is made under uncertainty. This means that those decisions are decision that have to be made under so-called inductive risk (cf. Rudner, 1953, Douglas, 2000, Steele, 2012). Inductive risk goes back to the problem of induction. The latter – in its original form – is the problem of justifying an inference from singular statements to a universal statement (i.e. a hypothesis or theory). The problem is that there is no deductively valid inference from the fact that all (so far) observed ravens have been black to the hypothesis: "All ravens are black". Accepting or rejecting a hypothesis like this requires a judgment as to when the evidence is good enough to make the (inductive) call or jump from evidence to hypothesis. Inductive risk is the "chance that one will be wrong in accepting (or rejecting) a hypothesis", i.e. the risk of rejecting a true hypothesis or accepting a false one. And the more uncertainty is involved in learning about something, the less (safe) evidence one has as a basis to make this decision, i.e. the higher the inductive risk becomes.

This kind of risk is not only present when accepting or rejecting hypothesis, but also in many other aspects of the scientific process (cf. Douglas, 2000). For example, in the domain of climate science, it is not only present when accepting or rejecting a hypothesis, but also much earlier, for example, when making assumptions about how to best represent the climate within a computer simulation. There is inductive risk about what the "true" structure of the climate is, how initial conditions should be assigned, how the current state of the climate can be best represented, and how the variables modelling the change of the climate should develop with time, just to name a few areas where inductive risk plays a role. The argument from inductive risk tells us that in such cases, there are multiple ways of making these decisions that are equally epistemically "correct" and well-founded, because of the gap there is between the evidence available to guide such

decisions and the evidence required to make the decision without risk.

As explained, there is much uncertainty, and many different kinds of uncertainty involved in learning about the climate. This means, that the problem of induction is especially prevalent here; and therefore the inductive risk associated with making decisions, here, is very high. And furthermore, climate scientists do not just rely one one single climate model to learn about the climate, but often, use many different models. Within each of those models, the modelling assumptions have to be made under inductive risk – for each model, there is no uniquely epistemically better or worse way of making these assumptions.

A consequence of this argument from inductive risk, and this will be important for the argument of this chapter, is that making different modelling choices (within different climate models) can lead to huge discrepancies between the hypotheses that are generated and justified by the resulting climate models.

It is important to note that this phenomenon does not make the models bad or not useful. Modelling the climate and making predictions about future climate change are, nevertheless, useful and informative. But sometimes, as we will see shortly, situations can arise that are difficult for orthodox epistemology to deal with. Climate science (and potentially other areas) requires an epistemology that can deal with this challenge posed by severe uncertainty. Fragmentation will turn out to be an attractive framework to deal with such situations.

The following is an example from the 2023 IPCC Assessment Report[2]. The IPCC is the so-called "Intergovernmental Panel on Climate Change" – a group of climate scientists that regularly produce reports on the latest developments within the earth's climate system[3]. Those reports are relevant

[2]Note that I erased probabilistic language from the example, since this account is qualitative. I am aware that climate models are probabilistic, but I think that the simplified qualitative version of looking at climate models is already epistemically fruitful. This project can and will be extended to a Bayesian version in the future.

[3]I am aware that the IPCC is a group consisting of multiple people. In this context, I consider them as an individual (group) agent, jointly producing and being accountable for the entirety of IPCC report.

for and used in both, climate science as well as climate policy.

The following claim[4], A, is taken from the IPCC report, and it expresses that "[g]lobal warming of 2 degrees Celcius will be exceeded during the 21st century" (cf. IPCC, 2023, p. 14[5]). Now consider something that might be surprising. The IPCC panel actually communicates not just A in their report, but also $\neg A$. They communicate that global warming of 2 degrees Celcius will be exceeded during the 21st century, but also that it will not be the case that global warming of 2 degrees Celcius will be exceeded during the 21st century, that is, $\neg A$ (cf. IPCC, 2023, p. 14[6]). The IPCC report, thus, contains both, the IPCC scientists' belief in A, as well as their belief in $\neg A$[7].

4.3 That's irrational! – Is it?

At this point, orthodox epistemology[8] would say, "wait a minute, there is something irrational going on here". One cannot both believe A and $\neg A$ at the same time, that would be irrational. But is this really a case of irrationality?

Let us look at whether and how the claims the IPCC makes in their

[4]I assume A to have the structure of a proposition, i.e. a set of possible worlds. This will be formally introduced in section 4.

[5]Citation:"Global warming of 2 degrees Celcius, relative to 1850–1900, [will] be exceeded during the 21st century (SSP3-7.0 and SSP5-8.5)."

[6]Citation: "Global warming of 2 degrees Celcius [will not] be exceeded (SSP1-1.9)."

[7]I assume the propositional attitude the IPCC scientists are communicating / asserting to be belief. If one were to argue for the relevant propositional attitude to be not belief, but merely something along the lines of acceptance, then the dilemma might still be there, but less intense. The reason for this is that, traditionally, the norms associated with belief are stronger than those associated with acceptance. I do not, however, consider this possibility here, because I think there to be good reason to believe that the IPCC scientists are communicating their beliefs. The main reason is the norms traditionally associated with assertion – for which acceptance is too weak of a propositional attitude to warrant assertion (cf. Pagin and Marsili, 2021).

[8]With "orthodox epistemology", I am referring to the idea that for an agent to be rational, their doxastic state needs to fulfil a set of strong rationality norms, such as overall consistency and closure under logic (cf. Borgoni, 2021), often referred to as the "unity model" or "web model" of belief. Stalnaker (1991) gives an overview over idealisations and normativity assumptions about belief.

report are justified. This will help to decide whether or not orthodox epistemology is right in classifying this as a case of irrationality, because if one of their beliefs were not justified, then clearly, it would not be good to assert them in the IPCC report. If this were the case, orthodox epistemology would have it easy, it would tell the scientists to get rid of the worse-justified claim. But as we will see, this is not the case here, and moreover, orthodox epistemology will not be of much help here.

At first glance, it seems that the IPCC scientists are not irrational. The reason for this is that their beliefs in A and $\neg A$ are well justified – they are supported by relevant evidence, and the scientists show that and how they are. The justificatory work is done by two respective and equally reliable climate models, on the basis of which the relevant claims can be predicted / synthesised. So, behind A and $\neg A$ stand two climate models, M and M', that support A and $\neg A$, respectively.

Both climate models, interestingly, are based on the same (historical) evidence, E, that expresses information about the development of the climate so far. The reason why, although both models are based on the same piece of evidence, E, they manage to justify jointly inconsistent claims, is that they operate on significantly different sets of background assumptions – which are equally epistemically warranted due to the severe uncertainty associated with learning about the climate. I will delve deeper into what exactly this means in a bit.

The scientists communicate their beliefs in A and $\neg A$, based on the respective justifying evidence, which is, the evidence E, relative to the respective climate models, M and M'. This is a problem for the orthodox epistemologist[9]. There can be no less justified claim given up if both claims

[9]I thank an anonymous reviewer for pointing out this potential challenge of interpreting the asserted beliefs in the IPCC report as being inconsistent. The challenge is that one could just interpret what they assert not as outright beliefs in A and a belief in $\neg A$, but as conditional beliefs of the form $B((M \wedge E) \rightarrow A)$, as well as $B((M' \wedge E) \rightarrow A')$.

The first reply to this is that, strictly speaking, the scientists do not assert conditional beliefs that take this logical form. The second point is that the appropriate conditionals would not be $B((M \wedge E) \rightarrow A)$, as well as $B((M' \wedge E) \rightarrow A')$, but instead, $B((M \wedge E) \rightarrow A)$, as well as $B((M' \wedge E) \rightarrow \neg A)$. $\neg A$ is much stronger than A', because A' is not necessarily inconsistent with A, but $\neg A$ is. Furthermore, A and $\neg A$ respectively are

are equally well-justified.

What are some other strategies that the orthodox epistemologist might offer to get out of this? They have two options. The first one is to adopt a cautious strategy and tell the scientists to suspend judgment on the matter of whether or not A is the case. But, the IPCC scientists might reply, why should I suspend judgment if I have such good justification for both A and $\neg A$? If this were the right strategy, the standard of justification for any belief would be too high to ever be attained in any complex system science, because what more should an IPCC scientist possess in favour of a claim than relevant evidence and a reliable and trustworthy climate model? So, this is not the right strategy here.

The second option the orthodox epistemologist might go for is to adopt a brave strategy and tell the IPCC scientists to apply doxastic logic (or belief revision theory[10]) to result in a consistent set of beliefs (by enforcing a closure requirement). In other words: "You should throw onto the epistemic trash either A or $\neg A$." But, the IPCC scientists might reply: "How should I throw either of these claims away if both of them are equally well-justified?" Both tell us something about future climate change that is relevant and justified [11] [12].

---

not, strictly speaking, (materially) implied by the conditionals, which is why I will use non-monotonic justification relations to represent the way in which these propositions relate to each other.

And lastly, it is not immediately clear how to believe a computer model. The contents of beliefs are usually taken to be propositional, but a climate model is a piece of code in a computer, so at least prima facie, a climate model does not have the right structure to be believed by the scientists. What can be believed are the inputs and outputs of the models, E, as well as A and $\neg A$. The models themselves can be thought to be reliable and trustworthy, but not believed by themselves. They, rather, provide sets of background assumptions relative to which the claims in question need to be considered. Without beliefs in M and M', the conditional beliefs would then take the form $E \rightarrow A$, and $E \rightarrow \neg A$, which would be equally or even more challenging for orthodox epistemology.

[10] Eg. Alchourrón, Gärdenfors, and Makinson (1985).

[11] Note that there is a third traditional strategy available that I gloss over in this chapter – going probabilistic. I do not discuss this strategy here, because this chapter deals with all-or nothing attitudes, and not probabilistic ones. Evaluating this potential choice (within a setting allowing for degrees of belief) would go beyond the scope of this chapter. In the following chapters, I will develop a probabilistic account of fragmentation.

[12] A proponent of either strategy might still insist that allowing inconsistencies among

This is the dilemma – we seem to have encountered a situation where a belief set of overall inconsistent claims seems rational, because those claims are well-justified. But orthodox epistemology tells us that this cannot be rational, and their toolbox for dealing with the situation (either the cautious or brave strategy) is unsatisfactory.

Fragmented justification will be a third and much more appropriate strategy that provides a more attractive alternative compared to the two extremes – the cautious and the brave strategy – that the orthodox epistemologist just offered. Fragmented justification will be a way to reframe the epistemic situation in a way that gets rid of outright inconsistency, as epistemic orthodoxy requires, while providing the reasons for why this inconsistency arose in the first place – the conflicting background assumptions the scientists have to make under severe uncertainty – and without having to give up on either of the well-justified claims and thereby being able to permit the preservation all of the justified claims, not just some of them. The fragmented justification solution will be superior to the at-

---

one's belief state is too costly. And, they might argue that it would be easy on either strategy, to not outright believe the claims that, either one suspends judgment on, or that are thrown onto the epistemic trash, but instead to merely "keep them in mind" so they can be still used for further inquiry. Neither the cautious nor the brave strategy would allow the scientists to believe both A and $\neg A$ in this scenario, but the attitude towards them (or one of them for the brave strategy) would have to be something weaker, such as "keeping in mind".

But this does not solve the problem. Belief can be considered to be a kind of "lower bound" for norms of assertion (usually the relevant norm for assertion is a knowledge-norm which is even stronger than belief, cf. Pagin and Marsili, 2021). This means that, agents should only assert claims they believe in. Claims an agent does not believe should not be asserted, according to this norm. Descriptively, the IPCC scientists do assert both A and $\neg A$ in their report – and seem to be doing so rationally, which both the cautious and the brave strategy cannot explain or deal with. A non-orthodox argument for why the IPCC scientists would be justified in asserting either claim if they do not hold a belief in that claim, or an argument against the intuition that those assertions seem rational given the good justification the claims enjoy would be required here from the orthodox epistemologist.

Another potential objection to this argument would be along the lines of Mandelbaum (2014), who argues that "keeping in mind", at first glance might seem like a propositional attitude different from belief, but that, in the end, any form of thinking is essentially believing. So, on this second reply, arguing that one can solve the problem by having the IPCC scientists merely "keep in mind" some of the claims brings back the problem from the back door, because "keeping in mind" can be considered a form of believing.

tempts at a solution that orthodox epistemology provides, because it possesses the resources to explain how the inconsistent conjunction of claims can be well-justified without having to require global consistency[13].

Let us take a closer look at A and $\neg A$ and the two climate models M and M' producing them. Climate model M, together with evidence E, justifies A. This is the case because of the assumptions that are made within model M. Model M assumes a high / intermediate emissions scenario (cf. IPCC, 2023, p. 14). This means that the variables and parameters of the model are set in such a way as to simulate a world where the emissions will be high / intermediately high for the considered timeframe. The assumptions that are thereby built into the model have to be made under uncertainty. It is not clear how *exactly* a high emissions world would look like, what *exactly* constitutes a "high emissions scenario", and how this can be best represented in the model. Furthermore, as mentioned already, the initial conditions of the model (as-is state of the climate) can only be determined with less-than perfect accuracy, because of the high complexity of the system.

And, there is meta-uncertainty about the model. Is it the right model to represent this scenario, or might there be a better one? Are all the assumptions within the model made in a way that it is sufficiently informative about the real climate? And, is the high / intermediate scenario the scenario that best represents the future state of the world? The future state of the world depends on decisions yet to be taken, so this is very difficult to epistemically evaluate at time of modelling.

[13]Note that fragmentation is not the only way of lowering the high rationality standards of orthodox epistemology. Other paths include weakening the logic that belief ought to adhere to (eg. by denying certain closure conditions, or consistency), weakening the concept of belief (and how it ought to cohere with probability) (cf. eg. Dorst, 2019), or introducing other contextualist notions of belief that diverge from fragmentation (Leitgeb, 2017). Another path is to reinterpret the rationality conditions of accounts of belief to be norms that don't necessarily have to be adhered to by rational agents (eg. the norms may be understood in terms of "constitutive norms" (cf. T. Williamson, 2000). However, the fragmentation strategy developed here is different in that it not only relaxes orthodox rationality requirements on belief, as could be done eg. by going non-classical, but instead, develops and uses an explanation for why the puzzle arises in the first place – the fragmented nature of the justification of the claims involved.

The other model, M', together with E, justifies $\neg A$, i.e. that global warming of 2 degrees Celcius will not be exceeded during the 21st century. The model M' is an equally reliable model as M. How can it be that it supports a claim that is inconsistent with A?

The reason for this, again, are decisions in the modelling process that have to be made under uncertainty, and where there is no purely epistemically better or worse way of making these decisions (i.e. the decisions have to be made under inductive risk). M' assumes a low emissions scenario (cf. IPCC, 2023, p. 14). This means that the parameters and variables in the model are set in a way to simulate a low emissions world. As before, the assumptions that are thereby built into the model have to be made under uncertainty. It is not entirely clear how *exactly* a low emissions world would look like, what *exactly* constitutes a "low emissions scenario", and how this can be best represented in the model. Furthermore, again, the initial conditions of the model (as-is state of the climate) can only be determined with less-than perfect accuracy. And, again, there is meta-uncertainty about the model, and the question whether the assumption of a low emissions future best represents the future state of the world. As before, the future state of the world depends on decisions yet to be taken, so this is very difficult to epistemically evaluate at time of modelling.

It is important to stress that both models agree on E – the claim expressing the development of the climate so far. There is no uncertainty involved here, it is known how, until the present moment, the climate has developed. Uncertainty comes in when trying to model and predict future climate change. But the fact that both A and $\neg A$ are based on the same piece of evidence E might appear as striking. The reason why the same piece of evidence can justify A within model M, and justify its negation $\neg A$ within M', is the uncertainty and inductive risk associated with making modelling assumptions and predicting future climate change. This uncertainty and inductive risk results in different but equally valid sets of background assumptions being built into the two different models.

Because of the complexity of the system and the uncertainty associated with learning about the system, as well as due to how climate modelling

works, situations like this can happen. But this does not mean that one should not trust the climate scientists anymore, because they are irrational. The predictions both models make are generated respecting the highest scientific standards and should both be taken seriously. The IPCC scientists should therefore not suspend judgment on either A or $\neg A$. They should neither use belief revision to get rid of either A or $\neg A$, because there is no reason to get rid of one over the other claim. These are the two options that the orthodox epistemologist would propose here.

Instead, the model that is most appropriate for such a situation is a fragmentation model of justification. This model will be able to reframe the epistemic situation by respecting the fact that both claims A and $\neg A$ are well justified (albeit generated under severe uncertainty), while getting rid of any outright contradiction.

The idea is that, corresponding to the two different models, there will be two fragments, F_1 and F_2. Within each fragment, different background assumptions and modelling choices are made, which leads to different plausibility assignments over claims, such as A and $\neg A$, in different fragments. Justification will be a relation between pieces of evidence and claims in need of support that is defined on fragments. Relative to the first fragment, A will be justified (but not $\neg A$), whereas relative to the second fragment, $\neg A$ will be justified (but not A). This way of framing the epistemic situation of the IPCC scientists makes space for them being rational in having justification for jointly inconsistent claims, while being able to get rid of the outright inconsistency, because A will be justified relative to one fragment, and $\neg A$ relative to another, but not both within the same fragment.

Note that this does not entail that fragmentation should be used whenever and wherever to escape any epistemically precarious situation by artificially rationalising what orthodox epistemology deems irrational. Fragmentation is an option for certain situations of epistemic pressure that seem rational but that orthodox epistemology deems irrational while not providing an epistemically satisfactory solution to them (like suspension of judgment or belief revision). Fragmentation is not a strategy that one can or should use in any kind of circumstance. Instead, it offers a way of

rectifying situations as rational that orthodox epistemology does not provide a sufficiently satisfactory analysis to. And actually, it appropriately describes what the IPCC scientists are actually doing in their report to rationalise the inconsistency. The IPCC scientists assert both A and $\neg A$, but contextualise both propositions. They explicate how A is justified relative to certain modelling assumptions (one fragment), whereas $\neg A$ is justified relative to other modelling assumptions (a different fragment).

Interestingly, when they synthesise their findings, the IPCC scientists also engage in a process of de-fragmentation – they extract what all fragments / models agree on[14]. De-fragmentation is a topic that has not yet been discussed in the fragmentation literature but becomes very important as soon as the fragmented belief set in question becomes relevant to practical decision making. I will delve into this topic at the end of this chapter.

4.4 Formalisation

After having presented the example, I now show how this situation can be formally represented within a framework of fragmented justification.

Let W be a non-empty set of possible worlds[15]. Let W be finite for reasons of simplicity. By a proposition, I mean a subset of the total set of possible worlds W. A proposition P is true at a world w just in case $w \in P$. Some propositions will take the role of a piece of evidence, E, that can justify another proposition. Some propositions will take the role of a proposition, A, that is in need of justification.

What has been described as a fragment in the example will be formally represented by a total plausibility pre-order $\leq$ on W. A total pre-order $\leq$

[14] "Global surface temperature will continue to increase until at least mid-century under all emissions scenarios considered. Global warming of 1.5 degrees Celcius and 2 degrees Celcius will be exceeded during the 21st century unless deep reductions in CO_2 and other greenhouse gas emissions occur in the coming decades" (cf. IPCC, 2023, p. 14).

[15] I use a possible worlds setup in my account, since it is a standard practice in modelling propositions and varieties of doxastic attitudes.

is a relation that is reflexive (For all $w \in W$: $w \leq w$), transitive (For all $u, v, w \in W$: If $u \leq v$ and $v \leq w$, then $u \leq w$) and total (For all $v, w \in W$: $v \leq w$ or $w \leq v$). This order ranks the worlds in W from most to least plausible.

Plausibility pre-orders are taken from two sources of inspiration. The first is the belief revision literature (cf. Alchourrón, Gärdenfors, and Makinson, 1985, Grove, 1988), where it has been shown that the classical belief revision operator, "*", can be represented via total plausibility pre-orders. The second source is the literature on non-monotonic reasoning surrounding Kraus, Lehmann, and Magidor (1990) and Lehmann and Magidor (1992). Both of these sources use this machinery to study belief, while I will adopt it to study fragmented justification.

Intuitively, a total pre-order over W can be thought of as the background assumptions and modelling choices that an agent is operating on in a specific circumstance. Given different sets background assumptions (i.e. within different fragments), different propositions can become more or less plausible.

Justification, then, is a non-monotonic[16] relation, $\mid\!\sim_{\leq}$, that is defined on a total pre-order, $\leq$. Given a total pre-order, $\leq$, over W, one can define conditional and unconditional justification in the following ways[17]:

1. **Conditional justification**: A proposition A is conditionally justified by a proposition E relative to a pre-order $\leq$ ($E \mid\!\sim_{\leq} A$) iff in all most plausible E-worlds from $\leq$, A is true, where the most plausible E-worlds are, among all worlds in which E is true, those E-worlds that are most plausible / ranked the highest on the pre-order $\leq$[18].

2. **Unconditional justification**: A proposition A is unconditionally

[16]It is not necessarily the case that, if $E \mid\!\sim_{\leq} A$, then $E, E' \mid\!\sim_{\leq} A$. I assume non-monotonicity to capture the assumption that justification is defeasible, i.e. that there can always appear a "new" piece of evidence that undermines a previously justified proposition.

[17]The justification relation formally resembles Lewisian counterfactuals (Lewis, 1973), but is entirely distinct from counterfactual reasoning, because it has a very different interpretation.

[18]I assume that the most plausible E-worlds exist relative to a pre-order $\leq$.

justified relative to a pre-order $\leq$ iff in all most plausible worlds from $\leq$, A is true.

The philosophical motivation behind this way of modelling justification is a very intuitive pre-theoretic notion about justification. Justification (of a proposition), here, is neither perceived to be just evidence (as in evidentialism), not just a reliable process (as in reliabilism), but a relation between the proposition in question, and evidence. As evidentialism prescribes, evidence is important for justification, but what makes justification justification is the relation in which that evidence stands to the proposition in need of justification, given some plausibility pre-order, i.e. set of background assumptions (of an agent). And this is captured by the notions of conditional and unconditional justification that make use of the plausibility pre-orders on which justification relations are defined[19].

Any pre-order-indexed justification relation, $\mid\!\sim_{\leq}$, is assumed to be closed under logic. This means that within pre-orders (fragments), traditional closure conditions hold. For example, if A is conditionally justified by E relative to $\leq$ and B is conditionally justified by E, relative to $\leq$, then the conjunction of $A \wedge B$ is conditionally justified by E relative to $\leq$. Further properties of so-called logics of justification can be found eg. in Kraus,

[19]The literature distinguishes between propositional and doxastic justification. Propositional justification is thought to be an agent-independent notion. It is about whether or not a proposition is justified independently of whether the agent actually believes that proposition on the basis of its justification. Doxastic justification, on the other hand, is about whether an agent justifiedly believes a proposition, i.e. whether a proposition that is actually believed by an agent is believed justifiedly / on appropriate grounds (cf. eg. Turri, 2010). Although the account developed in this paper has epistemic agents in mind, it can be interpreted as both, an account of purely propositional as well as an account of doxastic justification. Purely propositionally understood, the account gives a rule, given different plausibility pre-orders over the universe W, which propositions would be justified (by which pieces of evidence or unconditionally), relative to a respective pre-order. But, the perspective that is much more interesting for the example at hand is doxastic justification. Plausibility pre-orders are thought to be formal representations of the way an agent's background assumptions influence the way they evaluate propositions. So, plausibility pre-orders are thought to be part of an agent's doxastic state. And given this, the account tells us when and how a specific proposition can be justifiedly entertained by the agent, namely, if it is either conditionally justified through some evidence or unconditionally justified, and if the agent is also aware of the valid justification relation, relative to a certain set of background assumptions.

Lehmann, and Magidor (1990).

Given this formal apparatus, let us look at a formalised version of the example presented before to see how it works and why it works in this specific situation. I will first present the propositions involved in the example and then present how the two pre-orders corresponding to the two fragments look like. I then show what follows from these orderings in terms of what propositions are conditionally and unconditionally justified in each fragment.

A_0: Global warming will continue during the 21st century.

$\neg A_0$: Global warming will not continue during the 21st century.

A: Global warming of 2 degrees Celcius will be exceeded during the 21st century.

$\neg A$: Global warming of 2 degrees Celcius will not be exceeded during the 21st century.

E: Proposition expressing emission data development so far (until now).

The figures show the two pre-orders $\leq_1$ and $\leq_2$. The most plausible world in $\leq_1$ is $\{A_0, A, E\}$, whereas the most plausible world in $\leq_2$ is $\{A_0, \neg A, E\}$. Notice that the ranking of all other worlds other than the most plausible one on both orders is non-strict. This is one of the advantages of total pre-orders – they allow for uncertainty regarding which of those worlds are strictly more or less plausible than another world[20].

According to the first pre-order, the following holds:

1. A_0 is unconditionally justified relative to $\leq_1$, because A_0 is true in all most plausible worlds.

[20]The intuitive reasoning for the ordering of the less than most plausible worlds presented in $\leq_1$ goes as follows: It is more plausible that A_0 is true than that A_0 is false. It is more plausible that E is true than that E is false. But it is more true that $\neg A$ is true if $\neg E$ is true than that $\neg A$ is true if E is true. The intuitive reasoning for the ordering of the less than most plausible worlds presented in $\leq_2$ goes as follows: It is more plausible that A_0 is true than that A_0 is false. It is more plausible that E is true than that E is false. But it is more plausible that $\neg E$ is true if A is true than that A is true if E is true.

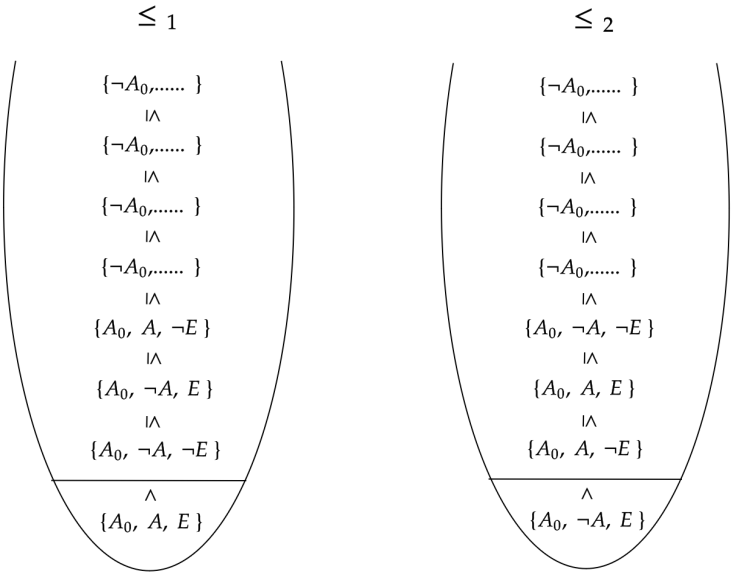

(a) Pre-order 1 (b) Pre-order 2

2. A is conditionally justified given E relative to $\leq_1$ ($E \mid\!\sim_{\leq_1} A$), because in all most plausible E-worlds, A is also true (namely in the single most plausible world)[21].

According to the second pre-order, the following holds:

1. A_0 is unconditionally justified relative to $\leq_2$, because A_0 is true in all most plausible worlds.

2. $\neg A$ is conditionally justified given E relative to $\leq_2$ ($E \mid\!\sim_{\leq_2} \neg A$), because in all most plausible worlds in which E is true, $\neg A$ is also true (namely in the single most plausible world).

[21]Note that maximal plausibility and maximal E-plausibility can diverge, although they coincide in this particular example. It could theoretically be the case that the most plausible E-world is not among the layer of most plausible worlds. Because they coincide here, strictly speaking, A_0 is also conditionally justified by E, and A is also unconditionally justified. But this does not change what I am trying to show with this example, as the first claim is weaker than the unconditional justification of A_0, and the second claim is stronger than the conditional justification of A which shows even more strongly the role that the difference in modelling assumptions between the two climate models plays. The same, but vice versa, will hold for the second pre-order.

Overall, A_0 is unconditionally justified relative to both orders / fragments. This makes sense, because both models predict that global warming will continue to increase. However, A is conditionally justified given E relative to $\leq_1$, and $\neg A$ is conditionally justified given E relative to $\leq_2$. In other words, A is only justified relative to model M, while $\neg A$ is only justified relative to M', interestingly, given the same piece of evidence on past emission developments, E. This shows that what makes the difference, here, is the different modelling assumptions the two different models make.

This way of representing the example is attractive, because it rephrases the epistemic situation in a way that does not require the agents to suspend judgment or do belief revision where at least one of the well-justified propositions would get lost. It represents how some propositions can only be justified relative to a certain pre-order, but not another, i.e. only relative to a certain set of background assumptions, but not another, but how some propositions may also be unconditionally justified by all pre-orders (i.e. how fragments can not only disagree but also agree on some things).

But why would an agent be rationally warranted to have two different plausibility pre-orders over the same set of possible worlds? The next section will deal with this question.

4.5 Reasons for fragmentation

Fragmented justification is not intended to merely be a descriptive account of how real agents cannot adhere to the strong rationality norms postulated by orthodox epistemology. The goal of this account is to be normative – in some situations, like the example above, fragmentation is the best epistemic option, it is what the agent should do in order to be rational. But this requires an argument beyond appealing to intuition that the agent is rational. And in many other situations, fragmentation will not be the best epistemic option, in those situations, there will be other options from the toolbox of orthodox epistemology to represent the situation that are less costly and better suited than fragmentation. In this section, I further

expand on the reasons for why justification can be rational in a situation like this, but not others.

The default doxastic state for an agent is unfragmented. If an agent believes that her Philosophy seminar is on Tuesday and also believes that her Philosophy seminar is not on Tuesday, at the same time, this does not warrant entering a state of fragmentation. What is happening here is that, either the agent does not have proper justification for the proposition that the seminar is on Tuesday or does not have proper justification for the fact that the seminar is not on Tuesday. But orthodox epistemology provides tools to deal with the situation. The agent should rethink what (she thinks) justifies her beliefs, and should ideally realise that one of them is not actually justified. In other words, the agent should do belief revision.

It is only when the agent finds themselves in a situation like the situation the IPCC scientists find themselves in, that becoming fragmented can become rational. I will call situations like these "situations of epistemic pressure". Situations like this are significantly different from the Philosophy seminar scenario, because here, both propositions A and $\neg A$ are equally well-justified and the options that orthodox epistemology offers for dealing with this situation do not respect this. And the reason for this is the uncertainty associated with learning about the climate system.

I will now delve deeper into how and why uncertainty, and related phenomena, can make it rational to fragment.

High levels of uncertainty can make fragmentation rational. It is not the case that the reason why the climate scientists assert inconsistent propositions is that they are making a mistake. After all, they are scientific experts, one would expect them to spot mistakes before writing an IPCC assessment report. Instead, something else is going on here; they have equally good justification for two jointly inconsistent propositions, A and $\neg A$.

And this happens, because the decisions and background assumptions they have to make when modelling the climate and predicting future climate change have to be made under inductive risk, which, to recall, is the risk of being wrong when making a decision under epistemic uncertainty.

Purely epistemically, there are multiple equally valid ways of making these assumptions, but different assumptions can generate very different predictions. So, given the uncertainty associated with the background assumptions one makes, and with the climate models themselves, it can be that different plausibility orderings over the same information (set of worlds W) are equally viable.

Let me go into a bit more detail here. Recall that the background assumptions that the scientists are making under uncertainty are different for model M that predicts A and for model M' that predicts $\neg A$. Model M predicting A makes assumptions that will make the model simulate a higher emission scenario than model M' predicting $\neg A$. Given that the models run simulations about the future, at the time of simulation, it is epistemically underdetermined whether or not the future world will look more like a low or high emission scenario, because this depends on non-epistemic (policy) decisions yet to be taken. Purely epistemically, the two models are equally good representations of different possible future states of the world – the possible future states differ from each other, but as of now, there is no way of deciding which state is more accurate or less. They generate inconsistent predictions, namely A and $\neg A$. But this does not mean that either model is faulty, that the scientists should suspend judgment completely, or that they should get rid of either A or $\neg A$. Instead, it means that A and $\neg A$ have come to be justified given different background assumptions. Given the assumption of a low emissions future, and the modelling choices that come with that, the proposition A that global warming of 2 degrees Celcius will be exceeded is much less plausible than given the assumption of a high emissions future and the associated modelling assumptions, and vice versa.

Note that, because modelling choices have to be made under uncertainty, one can even imagine a scenario where both models actually model a "low emissions scenario", but do so in different ways. There will be some modelling choices about how a low emissions model looks like that are underdetermined and for which, therefore, there is no unique right choice. Then, even in a situation where both models model the same category of scenario, radically different but equally justified propositions can come

about as well.

This is not to say that there is no right or wrong when it comes to modelling the climate, but that certain important modelling choices have to be made under inductive risk where there is no purely epistemic way of deciding which modelling choice is right or wrong.

This phenomenon is not a defect of rationality of the climate scientists, but a sign of a high level of complexity associated with the subject matter. Climate science is a complex system. Complexity and the uncertainty associated with learning about complex systems make is possible that a scientific expert can find herself being justified in incompatible propositions while knowing that only one of them can be true. The uncertainty surrounding climate models as well as their fragility with respect to modelling assumptions are responsible for the fact that such a situation can arise. Because there is such a high degree of uncertainty, having different and incompatible background assumptions in the background of the same set of propositions (manifesting in different and incompatible plausibility orderings over W) can be rational.

In cases where there are large amounts of uncertainty, it can happen that one finds oneself in a situation where one has good justification for jointly inconsistent propositions[22]. These situations of high uncertainty are situations where orthodox epistemology struggles to provide a solution, because it is unable to faithfully represent what is actually going on in such a situation. These are situations where fragmentation can make sense[23].

But note that this does not mean that in any situation where there is uncertainty involved, the agent is rationally warranted to fragment. Uncertainty is present everywhere in data-based sciences and everyday life. Many cases will not warrant fragmentation, because either one or the other proposition involved in the inconsistency is actually less justified than the

[22]Recall how in the previous chapter on the preface paradox, uncertainty was also a central reason for why fragmentation is rational in the context of that paradox.

[23]But, not every situation involving uncertainty will warrant fragmentation. Recall the Philosophy seminar example, where, of course, there is uncertainty involved, too. But in that case, albeit there being uncertainty, fragmentation is not warranted. In other words, uncertainty is not a sufficient condition to rationally warrant fragmentation.

other, or not justified at all (which makes it easy for orthodox epistemology to account for the situation). But there are rare cases where highly epistemically able agents like scientific experts come to be genuinely justified in jointly inconsistent propositions. And rather than getting rid of it by way of suspension of judgment or belief revision, like orthodox epistemology tells us, they should fragment, because this allows them to keep propositions that are perfectly well justified within their doxastic state which can be very useful for further inquiry.

Beyond uncertainty itself, there can be more factors closely associated with uncertainty, that can be reasons for why it can happen that an agent can have justification for jointly inconsistent propositions, and thus, have good reason to fragment. Let me present another one that plays a role here – differing but equally reliable sources of information.

A related reason for why it can be rational to have different plausibility orders over the total set of information can be if one works with different but equally reliable sources of information that provide incompatible background information, as it is the case in our example. Using two models, M and M', makes it possible in the first place to employ two different sets of background assumptions. Given different sources of information providing incompatible but equally reliable background information, different propositions will be more or less plausible to an agent. The plausibility evaluation of propositions can depend on the source(s) of information the agent is using, if they provide conflicting but equally reliable background information.

The climate scientist does not have a single source of information telling her to build both high and low level emissions assumptions into her models. Rather, she uses multiple different sources of information for her modelling purposes, some of which tell her about how a low emission future might look like, and some of which tell her about how a high emission future might look like. Different and equally reliable sources can contain inconsistent information due to uncertainty associated with the subject matter of the information. Given the uncertainty about whether we will actually live in a high or low emission future, and what exactly constitutes a high or low

emission future, different sources of information can be a another reason for justified but jointly inconsistent propositions.

One might think that, if the sources of information the agent uses are reliable or trustworthy, then differences in sources of information should not justify different plausibility evaluations and therefore different fragments (if one assumes the outside world not to contain inconsistency). But rather, one might think, the agent should incorporate both sources within one global doxastic state.

But again, if two sources are equally trustworthy, but contain incompatible information, the problem of the orthodox epistemologist outlined in the initial example tracks back to selection of background information. Why should an agent choose one source of information over the other if they are both equally reliable? Likewise, why should an agent not use either source, if they both are equally highly reliable (i.e. suspension of judgment on selection of background information)?

It can and it may happen that different sources of information, even if equally reliable, rationalise different plausibility orderings of the total set of worlds W.

This is the case for the climate scientist: The predictions of A and $\neg A$, respectively, are produced by two different but equally reliable climate models. Note that if the conflicting information, A and $\neg A$, were coming from one and the same climate model, then one would be more inclined to think that there must be something inherently wrong with the model, because it is inherently inconsistent. It is only because the conflicting information comes from two different but equally reliable models, that the situation can be rationalised.

As before, this does not mean that anytime an agent uses different sources of information for her reasoning, she is rationally entitled to fragment her information. In most cases, she should combine the information different sources provide in a consistent and unfragmented way. But if the agent finds herself in a situation where there are high levels of uncertainty involved, and where two (or more) equally reliable sources produce justified but jointly inconsistent propositions, and when orthodox episte-

mology cannot provide a solution to faithfully represent the uniqueness of this situation, then the agent can be entitled to fragment.

Large amounts of uncertainty, paired with scientific expertise, and the epistemic pressure of having equally well-justified but jointly inconsistent beliefs are the reason for why fragmentation is rational in a situation like this[24].

But one should employ fragmentation cautiously, that is, only when it is truly necessary; when orthodox epistemology is lost in appropriately accounting for a specific situation, like the one presented here. Introducing fragments complicates the doxastic state of the agent massively. It produces costs that are only worth paying if the potential benefits of fragmentation are big enough.

4.6 Precautionary norms for fragmentation

The last section concluded by mentioning that, if there are no reasons to go fragmented, one should not fragment, because then the costs of fragmentation would override the benefits. This leads to the topic of this section – two precautionary norms of fragmentation.

I will first present two general precautionary norms of fragmentation, and then motivate them.

1. Enter a state of fragmentation only when in a situation of epistemic pressure that cannot be sensibly resolved by the toolbox of orthodox epistemology.

2. If entering a state of fragmentation is warranted, only fragment as much as necessary to accommodate for the epistemic situation you find yourself in and respect the limits within and across fragments.

[24]There are more potential reasons for fragmentation, such as differences in stakes or subject matter. However, they are not relevant in this particular case, which is why I do not discuss them here, but some were discussed in the previous chapter on the preface paradox.

Fragmentation comes with benefits and it comes with costs. A benefit of fragmentation is that it allows an agent to keep all the well-justified propositions within their doxastic state, and not to have to suspend judgment or do belief revision. The benefit of fragmentation is that it provides a strategy for dealing with the epistemic situation the IPCC scientists find themselves in that is superior to the cautious and risk-seeking strategies of suspension of judgment and belief revision that orthodox epistemology provides. Is is superior in that it can explain and represent why the scientists are rational, albeit not adhering to the norms of traditional rationality, and furthermore, because it can solve the dilemma without making the scientists have to give up either well-justified belief.

Fragmentation allows to represent the epistemic situation in a new way that does justice to the specific epistemic conditions of the climate science example. It also provides a mid-way strategy in another related respect that concerns the normativity of classical logic for epistemic agents. Orthodox epistemology has it that rational belief should be conform with logic. This is, again, a cautious strategy. If an agent does not conform with the rules of logic (such as closure conditions, not allowing for inconsistencies etc.), they are irrational. On the other hand, there is a newer, more risk-seeking strategy about the rationality of agents – namely that classical logic is not the right normative guideline for the rationality of real agents. This position entails that logical closure conditions on the agent's doxastic state can be discarded. Again, fragmentation is a mid-way strategy between these two extremes. It does weaken the rationality criteria that orthodox epistemology imposes while not completely letting go of closure conditions.

Being able to frame situations like the IPCC climate science example in a new way is a significant advantage of fragmentation, because fragmentation offers an account that is less restrictive than orthodox epistemology, and thereby allows to classify highly epistemically able agents as rational (as they should be classified), despite them violating traditional rationality norms. At the same time, fragmentation needs to be restrictive enough to remain normatively significant, because fragmentation should not be an account of rationalising anything. I will shortly get in to these normative

restrictions on fragmentation.

But at the same time, fragmentation does come with costs. One such cost is that inferences can only be drawn within fragments, since any kind of logical closure properties only hold within fragments but not across fragments. The more fragmented an agent becomes, the more difficult it will be to draw inferences. Furthermore, each fragment will have its own fragment-relative justification relation. If there were very many fragments, one might worry that agents would not be cognitively able anymore to track all of these different justification relations and their differences, and thus, have significant limitations in their ability to draw meaningful inferences. A precise rule for how many fragments an agent can maximally rationally entertain cannot be formulated, since rational fragmentation is sensitive to multiple different criteria (as just presented) and the specific circumstance under consideration. But the general principle that should be followed is: "Introduce as many as necessary, but as few fragments as possible."

In the example case, this means that it would not make sense to fragment further beyond the two fragments that relate to the outputs of the two different models, respectively. Furthermore, since the fact that two models mean that two different but equally reliable sources of information existing is a reason for why it makes sense to fragment, it would not be obvious as to why one of those fragments should be fragmented further (i.e. to have multiple fragments stemming from eg. model M).

Beyond the general rule of only maximising as much as necessary to accommodate for the epistemic conditions of the situation at hand, there are some more specific norms one can formulate about the limits for fragmentation within, as well as across, fragments, specifics on which which the second precautionary rule depends.

4.6.1 Limits within fragments

How should fragments behave internally for a fragmented doxastic state to be rational?

The answer to the first question is partially given by the restrictions on

the plausibility orderings, the truth conditions of the justification relation, as well as by the closure conditions on the justification relation.

The plausibility orderings constituting the different fragments all have to be total pre-orders, that is, they have to be reflexive, transitive and total. Reflexivity means that for every world, that world is more plausible or equally as plausible as itself. This means that there can be equi-plausible worlds. Transitivity means that for all worlds, if u is more (or equally) plausible than v and v is more (or equally) plausible than w, then u is more (or equally) plausible than w. And totality means that for all worlds, either one world is more (or equally) plausible than the other, or the other way round. In other words, all pairs of worlds are comparable in that one precedes the other.

The reason why fragments are total and not partial is to represent the fact that, here, what distinguishes fragments is not which propositions are or are not accessible in what fragment (this would be accomplished by making the orders partial), but rather, the variability of background information that influences the plausibility assessment of the total information W (thus the order is total). If one chose to want to model the first kind, one could take fragments to be partial instead of total. The conditions on the plausibility orders just outlined already place some limits on how a fragment may permissibly look like.

Furthermore, recall the truth conditions of the (conditional) justification relation: "$(E) \mid\!\sim_{\leq} A$" is true iff in all most $\leq$-plausible worlds (in which E is true), A is true. This poses another limit on what is rationally permissible within a fragment and what is not. The relation only holds if the propositions in question are true in all most plausible (E-) worlds in that fragment. It would not be sufficient if they were true in some most plausible worlds, or, in worlds of any plausibility level. This makes it more difficult for the justification relation to be validated, i.e. for a certain proposition A to be justified.

Lastly, there are closure conditions imposed on the justification relation. Within a fragment, traditional closure conditions, such as closure under conjunction, and the explosion principle hold.

4.6.2 Limits across fragments

The intra-fragmentary norms of fragmented justification just discussed are in line with orthodox (unfragmented) epistemic rationality norms. But what about inter-fragmentary rationality norms, i.e. norms for how different fragments need to behave in relation to each other? Here, the present account significantly departs from the orthodox rationality requirements in that this account is relatively lenient with respect to inter-fragmentary rationality norms. This is to achieve the goal of the account to relax those orthodox rationality constraints for the account to be not too demanding for real agents to adhere to, while being demanding enough to remain normatively significant.

The main point about inter-fragmentary rationality norms is that different fragments, i.e. different plausibility orderings, may differ significantly from each other. Closure properties that hold within fragments do not have to hold across fragments. This means that one fragment might, in a very extreme case, assign maximal plausibility to a world w in which a proposition A is true, whereas a different fragment might assign minimal plausibility to world w, and instead, assigns maximal plausibility to a world v in which $\neg A$ is true. This leads to the possibility of a proposition A being justified in one fragment, but not in another, or, alternatively, two inconsistent propositions A and $\neg A$ being justified, each in a different fragment. Because closure properties do not hold across fragments, if A is justified relative to one fragment, and $\neg A$ relative to another, then it does not mean that $A \wedge \neg A$ is justified overall.

But note that this does not have to be the case; different fragments do not necessarily have to be incoherent with each other. But if they are coherent, then the precautionary norms presented above would demand checking whether the agent, indeed, finds themselves in a situation of epistemic pressure that cannot be solved by orthodox epistemology, otherwise fragmentation would not be warranted.

It is an asset of this account that it can rationalise doxastic states as the ones just described which would be called irrational by traditional accounts.

Still, for fragmentation to be able to be rational, surely, there must be some requirements on how different fragments should relate to each other.

There are two dimensions via which could take fragmentation to the extreme. The first one is increasing number of fragments. The second one is increasing the differences between the plausibility orderings involved. For both dimensions, one can impose limits to prevent extreme fragmentation on pain of irrationality.

For dimension number one: Fragmentation comes with benefits and costs and the more fragments are introduced, the more the costs will outweigh the benefits. One benefit of fragmentation is that it is able to rationalise cases of agents that traditional accounts label irrational but where there are strong intuitions that those agents are indeed rational, such as the climate science example. Fragmented justification is able to explain why agents can be rational in entertaining overall inconsistent sets of propositions, or why a certain proposition can be both justified and not justified at the same time for an agent. However, fragmentation comes with costs, as already mentioned, such as the fact that inferences can only be drawn within but not across fragments. The more fragments will be introduced, the more restricted the agent's capabilities of drawing inferences will be, and the more justification relations they will have to mentally keep track of. Therefore, as already explained, one should not introduce more fragments than necessary to solve the problem at hand, and thereby limit the number of fragments to avoid extreme fragmentation.

For dimension number two: One might want to have some limits on how different the fragments may look – especially if there are many. To do so, it is possible to stipulate a maximal distance measure between different fragments / plausibility orderings. It will depend on the specific scenario how exactly to choose such a measure – some scenarios will allow for more leniency regarding how different fragments and thus plausibility orderings may look for an agent to remain rational than others. Generally, an agent should never fragment more than is necessary.

4.7 De-Fragmentation

One might think that, while fragmentation might be epistemically warranted in some circumstances, one should worry about practical rationality when it comes to fragmentation. The reason one should worry is that, often, it will be difficult to rationally act upon a fragmented and overall inconsistent doxastic state, because it seems very difficult to find a rational course of action based on an doxastic state containing information such as A as well as $\neg A$. I mentioned that the IPCC reports are (or at least should be) relevant to designing and deciding on climate policies. Imagine a policy maker wanting to base climate policy on the basis of the IPCC scientists' doxastic state. How should this be done if that state is fragmented and overall inconsistent? This is why one should discuss options for getting rid of fragmentation to ensure practical rationality. This is a novel discussion within the fragmentation literature and the purpose of this last section.

The discussion of rational action and de-fragmentation is closely linked to the topic of how the account of fragmented justification relates to the concept of belief. In this section, I, first, argue that the account of fragmented justification is compatible with fragmentation accounts of belief. Furthermore, I will argue that the account of fragmented justification helps provide a rationality argument for account of fragmented belief. Since the question of rationality of fragmented belief is one of the most important questions in the literature on fragmented belief, one of the attractive features of the account of fragmented justification developed here is that it can boost the attractiveness of fragmented belief accounts by providing an additional rationality argument for them. I will keep the topic of the relationship between the account of fragmented justification and accounts of fragmented belief rather concise, since the main focus will be how to reach unfragmented belief, since this is what is commonly assumed to be the appropriate basis for rational action.

Proper de-fragmentation for ensuring practical rationality requires a suitable way to get from fragmented justification to unfragmented belief. I provide an overview over potential de-fragmentation mechanisms

and explain how finding a suitable way of getting rid of fragments is a highly non-trivial endeavour that will often not yield satisfactory results. I will argue that, often, no uniquely epistemically best aggregation mechanism exists, and that non-epistemic values will have to influence the de-fragmentation process, because otherwise de-fragmentation would often not yield any meaningful results.

4.7.1 Compatibility with fragmented belief

First, let us see why the account of fragmented justification is compatible with fragmented belief. A detailed discussion of this is outside the focus of this chapter, but the compatibility is important for further arguments I make in this section. I show that the common notion of belief used within plausibility frameworks similar to the one employed in the account of fragmented justification developed here has a fragmented sibling that is compatible with the present account of fragmented justification. There will be many more and more sophisticated versions of fragmented belief that can also be compatible with the present account, but working them out would go beyond the scope of this chapter, and would be an interesting avenue for further work, beyond this dissertation.

Within plausibility frameworks similar to the one that the account of fragmented justification makes use of, the most common notion of belief is that those propositions in the most plausible worlds are believed by the agent (Baltag and Smets, 2011, Baltag, Fiutek, and Smets, 2016, Artemov, Fitting, and Studer, 2024). For this notion, it is very easy to come up with its fragmented sibling. For any fragment / any plausibility pre-order, the beliefs relative to that fragment / pre-order will be the propositions within the worlds that are within the most plausible worlds according to that pre-order. This is the simplest and most closely related notion of fragmented belief that that is compatible with the account of fragmented justification developed here. They are compatible with each other, because one can define fragmented belief and justification, all within one and the same framework.

There are more possible and more sophisticated notions of fragmented belief that one could come up with which would be compatible with the present account of fragmented justification (which could include additional constraints like, eg. a total evidence constraint, cf. Carnap, 1947, Good, 1966), but this is outside the scope of this project.

4.7.2 Boosting the attractiveness of fragmented belief

Second, I will show how one of the attractive features of the account of fragmented justification is that it is able to boost the credibility of fragmented belief accounts by providing a rationality argument for them.

Unlike fragmented justification, fragmented belief is already an established research field and the question whether fragmented belief is or can be rational is being debated by some authors, such as Borgoni (2021) and Yalcin (2021). All of them argue that fragmented belief can, in fact, be rational, and thereby, that fragmentation of belief is not, per se, a failure of rationality. All of them also agree on the rationality norms within fragments, i.e. that belief fragments need to be internally consistent or coherent. Where they differ is what they deem the appropriate inter-fragmentary rationality norms to be.

The argument I will provide for the rationality of fragmented belief will be independent of, but will also strengthen rationality arguments such as Borgoni (2021) and Yalcin (2021) already present in the literature. I will first present my argument and then explain why it strengthens the arguments already present in the literature.

One of the main drivers behind the interest in the concept of justification is that whether a belief is rational or not is commonly thought to depend on whether that belief is well justified or not. What renders a belief rational is its justification. Now, imagine an agent that is fragmented with respect to belief. That agent entertains a number of belief fragments which each contains belief propositions. Some of the propositions in some fragment might be inconsistent with some propositions in a different fragment. If what renders a single belief rational is its justification, what will render a

fragmented set of beliefs rational must also have to do with the justificatory status of the beliefs contained in the different fragments.

Let me take an intermediary step – what renders an unfragmented set of beliefs rational? What renders an unfragmented and overall consistent set of beliefs rational is the respective justification that each belief within the set enjoys. Now if the structure of the belief set changes from an overall consistent and unfragmented structure to a potentially overall inconsistent and fragmented set, what can rationalise such a fragmented belief set, as well as the beliefs contained in the different fragments, is if the justifications supporting said beliefs also exhibit a fragmented structure and if this fragmented structure is warranted (as explained in section 5 and 6 of this chapter).

In short the argument is the following: What renders a fragmented belief state rational is the rational fragmentation of its underlying justifications. Of course, in a specific situation where an agent exhibits a fragmented belief set, it is not automatically given that her underlying justifications are fragmented in a rational way. This is something that needs to be checked in order to be able to assess whether there can be made a strong argument for the rationality of the fragmented belief state via the fragmentation of its underlying justifications.

How does this argument strengthen other rationality arguments for fragmented belief? Yalcin does not provide a positive argument but a number of intuition boosts using analogies between fragmented belief and rational disagreements of multiple agents or different (incompatible) perspectives in inquiry, where consistency or coherence is/should not be required across perspectives or inquiries.

Borgoni, like Yalcin, argues that fragments do not have to be consistent with each other to be rational, but poses an alternative inter-fragmentary rationality criterion: Responsiveness to evidence. The idea that, across fragments, beliefs must be responsive to the available evidence can be understood as follows: "While the subject only thinks with the belief fragment that is currently active in certain circumstances, the person's inactive beliefs stand as pieces of evidence to which that active fragment ought to be

sensitive" (Borgoni, 2021, p. 146). In other words, Borgoni argues that evidence, i.e. justification, is an unfragmented, global (and globally accessible) property of the agent's belief set that any given belief in any given fragment should be sensitive to in order to qualify as rational.

This means that beliefs in some fragment must be in some way sensitive to the justificatory evidence within that fragment, but also to justificatory evidence across fragments. This argument is not necessarily an argument for allowing overall incoherent or inconsistent belief states. Borgoni argues that an agent's inactive beliefs in other fragments serve as the role of evidence for the beliefs within the fragment that is currently active. However, she does not specify further what exactly the criterion of responsiveness to evidence entails. On one possible reading, it simply means that the beliefs in the agent's active fragment ought not contradict to the relevant evidence / beliefs the agent might harbour in other inactive fragments. But then, this would be nothing other than requiring inter-fragmentary consistency which seems, at least prima facie, to defeat the purpose of going fragmented.

If the reading should be something weaker, and the evidential beliefs in other fragments do not necessarily need to be consistent with the beliefs in the active fragment, it is not entirely clear what the "sensitivity" requirement stipulated by Borgoni means. It could mean something like: "Beliefs in other fragments that have a similar subject matter to the belief in question should be consistent with the belief in question". However, this would again boil down to a consistency requirement, restricted to propositions about the same subject matter.

Assuming justification to be fragment-relative significantly strengthens such rationality arguments for fragmented belief. The first argument by Yalcin (2021) is strengthened because relativising justification to fragments boosts the intuition that in fact, the fragment-relativity is so strong that there is no separate viewpoint from which inter-fragmentary inconsistencies might be perceived.

The second argument by Borgoni (2021) is strengthened because fragmented justification provides a concrete way on which beliefs ought to be "responsive to evidence" that, importantly, does not potentially boil down

to some version of requiring consistency across fragments: Beliefs in a certain fragment ought to be responsive to the evidence in the way that (belief) propositions ought to stand in a valid fragment-relative support relation to one or many justificatory propositions.

4.7.3 Reaching the goal of unfragmented belief?

After having established that the account of fragmented justification is compatible with fragmented belief accounts, and that, in fact, supporters of fragmented belief accounts should adopt a supplementary account of fragmented justification to strengthen the rationality argument for fragmented belief, it is time to explore whether and how one can generate unfragmented belief starting from fragmented justification.

It is known from social choice theory, in particular, from Arrow (1950)'s impossibility theorem, that meaningful aggregation of preferences is not possible given a number of prima-facie plausible rationality postulates. In the judgment aggregation literature (cf. Dietrich and List, 2016) the same has been shown for the aggregation of propositions. An instance of that general result holds for plausibility pre-orders – some assumptions on what a desirable aggregation mechanism should possess will need to be given up in order to result in a unified overall plausibility pre-order (cf. eg. Rott, 2001, Bonanno, Delgrande, and Rott, 2012). Based on this seminal work, it will not be surprising that finding aggregation mechanisms for fragments that satisfy desirable properties, and, moreover, even generate useful / interesting results is not trivial.

There are two main routes one can take to de-fragment. Though, the first route is only a pseudo-route, because it does not actually get rid of fragments, but produces a set of unfragmented beliefs while keeping fragments intact. This route consists in aggregating on the sets of those propositions that are justified in each fragment, individually. Pre-orders / fragments remain present here (they are not being eliminate through the aggregation process), which is why I focus on the second route that genuinely aims at getting rid of fragmentation. It consists in aggregating pre-orders /

fragments into one single pre-order which will then represent the overall doxastic state of the agent[25].

There are many different specific options for aggregating pre-orders. Which aggregation mechanism to get rid of fragmentation will be the most suitable will highly depend on the specific plausibility orderings at hand. And it is important to note that, it will not always be possible to intelligibly merge different fragments to one unique plausibility pre-order. I will go though some of the most common options and explain how they might not be producing an intelligible result for certain types of plausibility orderings.

First, the union of fragments. The first problem with this option is that it is not clear what the union(-order) of two pre-orders supposed to be. Attempting to take the union of two plausibility orderings $\leq_1$ and $\leq_2$ would result in an overall plausibility ordering that contains both, $\leq_1$ and $\leq_2$. If $\leq_1 = \leq_2$, then the union would just be $\leq_{total} = \leq_1 = \leq_2$. As soon as $\leq_1$ and $\leq_2$ disagree on the plausibility of certain worlds, the overall plausibility ordering will not be able to satisfy its rationality criteria – because it would have to but cannot assign a world w two different plausibility rankings at the same time. Taking this to the extreme, if $\leq_1$ and $\leq_2$ were complete opposites of each other, then taking the union of $\leq_1$ and $\leq_2$ would not result in there being any interesting plausibility ordering anymore[26], or else, an incoherent ordering, which would be a very uninformative outcome of the aggregation process, especially with regards to the goal of rational practical action. The reason for this is that the more the orderings disagree on, the more worlds it cannot intelligibly order in the overall ranking, because $\leq_1$ and $\leq_2$ disagree about their ranking. One may only hope to receive an intelligible result when trying to take a union-order of two pre-orders.

Second, the intersection of fragments: Unlike for the union, the intersection of two pre-orders is a defined notion, but the resulting order will only be partial. Taking the intersection of two plausibility orderings results in

[25]The literature sometimes calls the distinction between these two routes "compromise vs consensus".

[26]For the interested reader: Let $\leq_1$ be $t \leq u \leq v \leq w$, and let $\leq_2$ be $w \leq v \leq u \leq t$. Then, the union would result in $t = u = v = w$.

an overall plausibility ordering containing only what both plausibility orderings "agree on". The more $\leq_1$ and $\leq_2$ "agree on", the more informative will the overall ranking $\leq_{total}$ be in the sense that it provides a plausibility rank for more and more worlds, the more worlds $\leq_1$ and $\leq_2$ agree on. But then, in a sense, the more $\leq_1$ and $\leq_2$ agree, the less informative the overall ranking becomes, because it becomes more and more similar to $\leq_1$ and $\leq_2$ themselves. Now, the less $\leq_1$ and $\leq_2$ agree on, the less informative the overall ranking will be in the sense that it will have to omit more and more worlds and do not assign them any ranking at all, since $\leq_1$ and $\leq_2$ disagree about their ranking. Similar to the union, this can be taken to the extreme with to exactly opposing orderings $\leq_1$ and $\leq_2$. Here, there would not be any overall ranking anymore[27].

Third, one might think that unanimity or majority aggregation procedures might be an attractive option. In the case of more than two fragments, majority could be an option. However, for both majority and unanimity, but especially unanimity rules, the difficulty with respect to fragmentation is that one of the key advantages of this account is that fragments may differ significantly from each other. And this makes such rules very difficult to be usefully applicable.

Another option, one might think, could be to compute the mean plausibility value for each world[28]. However, this option does not work for pre-orders, since there are no "plausibility values" attached to the different worlds in the ordering. One might think that, one could still introduce values through coding orders through numbers that one attaches to the worlds in the ordering. But again, those numbers are merely ordinal, and taking an average or mean is not defined for an ordinal sequence like a pre-order.

The less extremely opposed fragments are, the more options for viable aggregation mechanisms open up. But generally, it can be the case (as with completely opposed fragments), that there might not be any suitable ag-

[27]For the interested reader: Let $\leq_1$ be $t \leq u \leq v \leq w$,and let $\leq_2$ be $w \leq v \leq u \leq t$. Then, the intersection would result in each t, u, v and w being equally plausible with itself, respectively, but there not being any overall plausibility ranking between t, u, v and w.

[28]This could be interpreted as a variant of the Borda count (cf. eg. Pacuit, 2019).

gregation procedure available. In any case, choosing a suitable mechanism when in need of de-fragmentation is highly non-trivial.

What should one do about this? In some cases, de-fragmentation might never become necessary and so the question for how to aggregate does not arise. The first reason why this can be the case is that, fragments containing inconsistent information might not get activated at the same time for a certain practical decision. It might also be that they do get activated at the same time but where this does not pose a threat to practical rationality[29].

But what if de-fragmentation is indeed necessary to be able to ensure practical rationality? What about the climate science example? Imagine a policy maker who wants to decide what policy to implement on the basis of the newest justified scientific hypotheses. They will have significant trouble deciding on a policy if they receive the information that A, as well as the information that $\neg A$ from the IPCC scientists. How should the two plausibility pre-orders be aggregated?

One could decide to aggregate by extracting what both orders agree on, which is A_0, the proposition expressing the fact that global warming will continue during the 21st century. But this is a very uninformative overall doxastic state. It is not sufficiently precise to help the policy maker decide on any policies on how to best prevent, mitigate or adapt to future climate change. The policy maker needs to know whether global warming will exceed 2 degrees Celcius or not during the 21st century. The climate scientist has a problem here, because they have good justification for both. As just mentioned, using the intersection method, neither A nor $\neg A$ will remain in the de-fragmented doxastic state. Using the union, both would stay, which does not help the policy maker either. Using majority / unanimity rules will also not yield any interesting results, because, precisely, there are two fragments that disagree about whether A or $\neg A$ is justified. Using a

[29]Discussing this possibility goes beyond the scope of this chapter, but one example case of this is the preface paradox, where the author of a book seems rational in believing an overall inconsistent set of propositions. Here, contrary to the climate science case discussed in this chapter, this does not pose a threat to practical rationality.

mean plausibility measure might seem possible, but is impossible for the plausibility pre-orders under consideration here.

In such a case, my argument is that non-epistemic value judgments need to be employed to come up with a suitable aggregation mechanism. The argument for this is a variant (and new application) of the seminal argument by Rudner (1953) (with variants by Hempel, 1965, Douglas, 2000 and Steele, 2012). There is a multitude of aggregation options available but there is no purely epistemic "fact of the matter" as to which aggregation mechanism will be best given a specific de-fragmentation problem. Similar to the decision of background and modelling assumptions, deciding on a de-fragmentation mechanism is a decision that is underdetermined by evidence, and thereby, a decision that has to be made under inductive risk. Rudner (1953) argues that the "gap" between the evidence and the decision is to be filled by ethically evaluating how serious a mistake in making a false decision would be, i.e. by employing non-epistemic value judgments.

The argument for why non-epistemic value judgments should be employed when aggregating fragments is similar. When faced with a practical decision problem for which multiple conflicting fragments are relevant, and being unable to make a decision on purely epistemic grounds, the selection of an aggregation procedure should be guided by evaluations of minimising negative consequences when making a mistake.

In the case of the climate scientist, they should result in an overall set of "beliefs" that is consistent and that contains from A and $\neg A$ that proposition that minimises negative practical consequences if it were false. This means that the scientist should choose to accept A instead of $\neg A$, because falsely accepting that climate change won't be as severe as it will actually be, will cause less strict policy measures, which, in turn, can lead to far worse consequences than falsely accepting that climate change will be more severe than it will actually be.

Note how the term "belief" was just used in quotation marks. On the basis of these considerations, one can conclude that whatever results from either aggregation process should potentially not be called belief, but rather something like the set of "accepted propositions". The reason for this is

that through using an aggregation mechanism and employing non-epistemic values, the central goal of belief of "aiming at the truth" might get lost. Instead, the purpose of the "beliefs" generated is not necessarily to capture the truth as best as possible, but, rather, to be the best possible disposition for rational practical action.

These considerations only become relevant for practical action. Only when, after reading the IPCC report, the policy maker gets back to the scientists, and tells them that they cannot rationally act upon the asserted fragmented belief set, then the scientists should de-fragment. Purely epistemically, there is nothing wrong with having incompatible fragments. It is only when considering practical rationality that de-fragmentation by way of employing non-epistemic value judgments can become necessary.

4.8 Conclusion

The aim of this chapter was to present and motivate an account of fragmented justification. There are epistemic situations where the strategies of orthodox epistemology are unfitting tools. There are cases, like the IPCC climate scientist case presented and developed in this chapter, where agents seem justified in believing jointly inconsistent sets of propositions, and where orthodox epistemology can only provide unsatisfactory solutions to the problem this poses. The reason why this can happen is the fragmented nature of justification in situations of severe uncertainty. And orthodox epistemology cannot represent or explain this, while the account developed in this chapter can.

Furthermore, I have presented and elaborated on reasons for fragmentation, as well as some precautionary norms that should govern fragmented justification. I have explained how the account of fragmented justification is compatible with fragmented accounts of belief, and that, supporters of fragmented belief accounts should also adopt a fragmented account of justification, because fragmented justification can be viewed as a rationality argument for the fragmentation of belief.

I have shown how there is no generally applicable rule on how one should go about de-fragmenting one's doxastic state to result in unfragmented belief, when facing problems concerning practical rationality and action. I have argued, drawing from the seminal Rudnerian inductive risk argument, that the choice as to how one should de-fragment in a specific situation where fragmentation poses a threat to practical rationality needs to be guided by non-epistemic value considerations.

Subject to further work in the following chapters will be the development of a Bayesian / probabilistic version of fragmented justification. And furthermore, another avenue for fruitful further research that goes beyond the scope of this dissertation would be, as mentioned to investigate in more detail the varieties of accounts of fragmented belief (and their formal subtleties) that could be compatible with fragmented justification, and to come up with a unified account of fragmented justification and justified belief.

Chapter 5

Fragmented Bayesianism

This chapter moves from a qualitative account of fragmented justification to a quantitative or probabilistic one. It presents a very general version of a fragmented account of Bayesianism as an epistemic theory of justification. First, it will recall the central theses of orthodox Bayesianism, and provides an overview over the notions of justification that can be formulated within orthodox Bayesianism. It, then, presents a very general account of fragmented Bayesianism, and then provides a list of parameters along which the account can be made moderately or very fragmented, depending on which and how many of those parameters one chooses to relativise to fragments. It, furthermore, presents an argument for fragmented Bayesianism. I first provide a very general argument for fragmentation, and the provide arguments in favour of the fragmentation of each of the parameters. Lastly, the chapter concludes by a brief comparison between the account of fragmented Bayesianism developed here and the qualitative account of fragmented justification developed in the preceding chapters.

5.1 Introduction

Fragmentation accounts of doxastic attitudes have been explored extensively for binary (all or nothing) attitudes, such as belief (cf. Davidson, 1982b, Lewis, 1982, Stalnaker, 1984, and Borgoni, Kindermann, and Onofri,

2021). In previous chapters, I have developed a binary / qualitative fragmented account of justification (of binary belief). However, there is less work on a fragmented account of degrees of belief, let alone fragmented versions of quantitative accounts of justification.

In this chapter, I will develop a fragmented version of Bayesian justification, departing from some work that has already been done on fragmented versions of Bayesianism by Elga and Rayo (2022), Fleisher (2023), as well as the literature on imprecise probabilities (cf. Kyburg, 1968, Levi, 1974).

Orthodox Bayesianism provides multiple different notions of justification. I will cash them out and will provide fragmented Bayesianism as a possible way of dealing with the availability of multiple different and at times incompatible notions of justification present in orthodox Bayesianism. I present the account of fragmented Bayesianism, and I show how fragmented Bayesianism has multiple parameters along which it can be made more or less fragmented, and thereby, more or less close to traditional Bayesianism. I will argue, similar to how I did for fragmented qualitative justification, that the necessity to reason and make decisions under uncertainty is the main justifying factor for fragmented Bayesianism. Finally, I will compare the quantitative account of fragmented justification developed here to the qualitative account of fragmented justification developed in earlier chapters.

The chapter will be structured as follows: Section 2 recalls the central theses of orthodox Bayesianism, and provides an overview over the notions of justification that can be formulated within orthodox Bayesianism. Section 3 presents a very general account of fragmented Bayesianism, and then provides a list of parameters along which the account can be made moderately or very fragmented, depending on which and how many of those parameters one chooses to relativise to fragments. In section 4, I present an inductive risk argument in order to justify fragmentation. I first provide a very general version of it, and then provide (inductive risk) arguments for the fragmentation of each of the parameters. Lastly, the chapter concludes by a brief comparison between the account of fragmented Bayesianism developed here and the qualitative account of fragmented justification

developed earlier.

5.2 Orthodox Bayesianism

The slogan of orthodox Bayesianism can be summarised as this: Keep your degrees of belief (i.e. subjective probability / credence function) coherent (i.e. adhere to the axioms of probability) at all times, and update them by way of Bayesian conditionalisation as new information comes in.

The core of orthodox Bayesianism thus consists of two theses, probabilism and conditionalisation (the formulation of those theses I use roughly follows Vallinder, 2018). Let W be a finite (for reasons of simplicity) and non-empty set of possible worlds. Then $\mathcal{A}$ is a powerset algebra on W – it is the set of all subsets or powerset of W. This algebra comprises all propositions P within W. Then, Bayesianism is the combination of the following two theses:

1. Probabilism: A rational agent's degrees of belief are represented by a credence function c which is a function from $\mathcal{A}$ into the interval $[0, 1]$. This means that it assigns a credence value between zero and one to all propositions in the logical space. The credence function c respects the axioms of probability theory (Kolmogorov axioms):

   1) $c(P) \geq 0$, for all P

   2) $c(W) = 1$

   3) For incompatible (mutually exclusive) propositions P and Q: $c(P \cup Q) = c(P) + c(Q)$

2. Conditionalisation: A rational agent updates her degrees of belief over time by conditionalising her credence function on all the evidence she has received. If an agent with credence function c_0 learns new evidence E between t_0 and t_1, then her new credence function c_1 is given as $c_1(P) = c_0(P \mid E)$[1].

[1]Using the Bayes Theorem, this conditional probability can be reformulated which makes it easier to calculate in practice.

So far, this definition is similar to the definition of Bayesianism provided in the literature overview chapter[2]. However, there are some more parameters that are implicitly present in orthodox Bayesianism and that have to be made explicit to show how they react to the introduction of fragmentation into the account.

Firstly, a credence function comes with a notion of background knowledge, K, implicitly encoded into the function. The background knowledge, K, associated with a credence function c, is the set of worlds in W / propositions in $\mathcal{A}$, that receive non-zero credence by the credence function, c, i.e. $K = \{P_1, P_2, ..., P_n\}$, where $c(P_1) \geq 0$ and ,..., and $c(P_n) \geq 0$. In other worlds, the background knowledge K associated with a particular credence function c are those propositions that are epistemically possible according to that particular credence function. Conditionalisation always happens on the new piece of evidence received, but moreover, happens always implicitly on the information the agent has already gathered hitherto, i.e. their background knowledge K.

Secondly, a credence function c not only comes with a set of background knowledge, K, but it implicitly comes with something called an ur-prior, U, which describes the belief state of the agent "before they learn anything at all" (Fleisher, 2023). This means, the ur-prior describes the agent's belief state, before they have any kind of background knowledge K to use as their prior when learning new information and engaging in the process of conditionalisation. The ur-prior contains the agent's basic epistemic standards, and is used to conditionalise on while the agent does not have required any background knowledge yet, but is also always implicitly conditionalised on, even when the agent has background knowledge, because the agent's epistemic standards need to be "remembered" through her learning process.

Thirdly, as is already explicit in the definition, a credence function is defined on a specific algebra $\mathcal{A}$ which, in turn, is defined on a set of worlds

[2]In the introductory chapter, I have also introduced the principal principle as a cornerstone of Bayesianism which I am consciously omitting here, because it is less relevant with regards to fragmentation. The reason why it is less relevant is that if an agent knows the objective chance of an event, then there is not enough uncertainty involved to warrant fragmentation in the first place.

W.

And lastly, and very intensely discussed in the Bayesianism literature, is the notion of justification / the confirmation measure that is chosen to go along with a credence function c:

5.2.1 Notions of justification therein

Let us look at the different notions of justification that are present in orthodox Bayesianism and how they relate to the landscape of traditional accounts of justification. First, note that, generally, Bayesianism can be viewed both as a(n epistemic) theory of justification as well as as a decision theory (a theory of practical rationality), or both. In this chapter, I focus on Bayesianism as an epistemic theory of justification.

Bayesianism, understood as an epistemic theory of justification, contains multiple implicit notions of the concept of justification. They are implicit, because standard definitions of Bayesianism, by themselves, do not state any explicit definition of the concept of justification. Within Bayesianism, the concept of justification is there, but it is "hidden" behind and to be spelled out via the concept of credence. Commonly, what orthodox epistemology understands as "justification" is understood in Bayesianism as a confirmation or support measure (cf. eg. Fitelson, 1999, Hartmann and Sprenger, 2010, Crupi, 2021). Bayesianism provides an array of notions of justification, which I will label M, that are "hidden" behind the credence formalism:

a. absolute notions of justification

   i. Proposition E is a justification for proposition P (given background K, and given an ur-prior probability measure U) just in case $c(P \mid E \wedge K) > c(P \mid K)$

   ii. Proposition E is a justification for proposition P (given background K, and given an ur-prior probability measure U) just in case $c(P \mid E \wedge K) > t$, where t is a (Lockean) threshold that may be set, eg., to 0.5 or 0.9 etc.

iii. combinations of i. and ii.

b. ordinal and quantitative notions of justification

iv. Proposition E is a stronger justification for P than proposition E' (given background K, and given an ur-prior probability measure U) just in case $c(P \mid E \wedge K) > c(P \mid E' \wedge K)$

v. The degree of justification of proposition P is x just in case $c(P) = c(P \mid E \wedge K) = x$

c. mixed notions

vi. combinations of i.-v.

Within the literature on Bayesian confirmation or support measures that can be interpreted as different notions of justification, it has been shown that a lot hinges on the choice of the concept of justification / confirmation measure[3]. The validity of important arguments within Bayesianism hinges on the choice of confirmation measure (eg. shown via the so-called Popper-Miller argument, cf. Popper and D. Miller, 1983, Gillies, 1986, Fitelson, 1999[4].). These different notions of justification are not necessarily compatible with each other and can yield wildly different results about the kind of claims that will be well-justified / well-confirmed.

Importantly, and this makes the topic of choice of concept of justification so widely discussed, is that, purely epistemically, there is no uniquely better or worse concept of justification / confirmation measure. They are

[3]In the confirmation measure literature, some of the most prominent confirmation measures go under the following names: difference measure, counterfactual difference measure, the log-ratio measure, the log-likelihood measure, Carnap's relevance measure, cf. eg. Fitelson (1999), Hartmann and Sprenger (2010) and many others (cf. eg. Tentori et al., 2007, Atkinson, 2012, Fitelson, 2021). In fact, those confirmation measures fall in the category non-absolute / relative notions of justification above. But I will not discuss them in detail, because the more abstract classification suffices for the purposes of this chapter.

[4]The argument fundamentally attacks Bayesianism as a theory of inductive probability and argues that "all probabilistic support is purely deductive". However, the argument hinges on a property of a specific notion of justification / confirmation measure – the additivity property of the so-called difference confirmation measure.

able to achieve different things, and can be helpful for different purposes. Therefore, choosing a suitable concept of justification to employ is a choice that has to be made under inductive risk. Inductive risk is the risk of being wrong when making a choice under uncertainty.

There are multiple proposals in the literature on how to deal with making decisions under inductive risk, such as Rudner (1953), Jeffrey (1956), Douglas (2000), and Steele (2012). In his seminal argument, Rudner argues that in order to make decisions under inductive risk, i.e. to bridge the gap between the available evidence and a decision, one needs to employ non-epistemic value judgments. Douglas (2000) and Steele (2012) extend this argument to various other parts of the scientific process, as well as to the intersection between science and policy. Jeffrey (1956) takes another route, and argues that it is not necessary to employ non-epistemic value judgments, despite the presence of inductive risk. Instead, he argues that, when dealing with a particular topic, such as the effectiveness of a certain vaccine, the scientist's job is not to finally accept or reject a hypothesis about the topic, i.e. it is not their job to make a decision under inductive risk. Instead they should simply stop when they arrive at a probability distribution on the basis of their research. Like this, they can avoid employing non-epistemic values – by not bridging the gap from evidence to the acceptance or rejection of a hypothesis at all. However, to simply come up with a probability distribution over outcomes, certain decisions under inductive risk need to be made, too. One of them is the choice of confirmation measure, which is underdetermined by evidence.

Fragmentation will provide another way of dealing with the inductive risk associated with certain decisions one has to make when learning about the world and updating one's doxastic state, such as the choice of an appropriate confirmation measure.

After introducing the parameters of background knowledge K, ur-prior U, algebra $\mathcal{A}$, and notions of justification / confirmation measure M, let me state again a definition of orthodox Bayesianism in a way that contains

these parameters[5]:

Let W be a finite (for reasons of simplicity) and non-empty set of possible worlds. Then $\mathcal{A}$ is an algebra on W; it is the set of all subsets or powerset of W, and it comprises all propositions P within W. Then, Bayesianism is the combination of the following two theses:

1. Probabilism: A rational agent's degrees of belief are represented by a credence function c which is a function from $\mathcal{A}$ into the interval $[0, 1]$. This means that it assigns a credence value between zero and one to all propositions in the logical space. The credence function c respects the axioms of probability theory (Kolmogorov axioms):

   1) $c(P) \geq 0$, for all P

   2) $c(W) = 1$

   3) For incompatible (mutually exclusive) propositions P and Q: $c(P \cup Q) = c(P) + c(Q)$

2. Conditionalisation: A rational agent updates her degrees of belief over time by conditionalising her credence function on all the evidence she has received. Let U, the ur-prior for c, encode the agent's epistemic standards. Let K be the background knowledge the agent possesses, i.e. the set of live possibilities, the propositions that receive non-zero probability. If an agent with credence function c_0 learns new evidence E between t_0 and t_1, then her new credence function c_1 is given as $c_1(P) = c_0(P \mid E \wedge K\ (\wedge U\ ))$[6].

[5]The only parameter missing from this definition is the notion of justification / confirmation measure, simply, because this parameter is not set when defining orthodox Bayesianism. In other words, the definition orthodox Bayesianism is compatible with / allows different notions of justification within the account.

[6]Using the Bayes Theorem, this conditional probability can be reformulated which makes it easier to calculate in practice.

5.3 Fragmented Bayesianism

This section introduces an account of fragmented Bayesianism. There is some work on fragmented versions of Bayesianism already present in the literature, however, it is not much, and this account will be more general than the ones postulated so far. Depending on how one chooses to set the fragmentation parameters I will introduce, and some additional assumptions, the accounts already existing by Elga and Rayo (2022), Fleisher (2023) will be variants of this account. But the aim here is to provide a very general formulation of fragmented Bayesianism and to spell out the dimensions along which one can make the account more or less fragmented, and thus, more or less close to orthodox Bayesianism.

According to fragmented Bayesianism, an agent's doxastic state can be represented not by one, but by a set of Bayesian credence functions, $c_1, ..., c_n$[7]. These credence functions, $c_1, ..., c_n$, will, either come with their own algebra $\mathcal{A}$, set of background knowledge K, ur-prior U and notion of justification / confirmation measure M, or all of them or some of them can share the same algebra $\mathcal{A}$, set of background knowledge K, ur-prior U and notion of justification / confirmation measure M.

Each credence function corresponds to a mental fragment of the agent. Each fragment, internally, ought to adhere to the rules of orthodox Bayesianism. This means, each individual credence function ought to obey to the axioms of probability, and updating should happen, in each fragment internally, via the conditionalisation process. There will be multiple options for how this conditionalisation process in a fragmented setting can look like exactly. I set fragmented conditionalisation aside in this chapter, but will talk more about it in the subsequent one.

What makes fragmented Bayesianism significantly different from orthodox Bayesianism is that not one but multiple credence functions are allowed within one single doxastic state. And with that comes that it suddenly becomes permissible for different fragments to have assigned different levels

[7]I assume a finite set for reasons of simplicity.

of credence to potentially one and the same proposition. Fragmentation allows for doxastic states containing epistemically conflicting information. Of course, there is need for an argument for why being in such an epistemic state could be rational – I will give en inductive risk argument for the rationality of fragmented Bayesianism in section 4. Before that, I will now delve into a list of (formal) parameters that the account of fragmented Bayesianism contains, and that can be adjusted to render the account more fragmented, or closer to traditional Bayesianism. They can also be adjusted to model different types of fragmentation.

5.3.1 Parameters

I will now introduce a number of parameters into the basic formal structure of fragmented Bayesianism consisting of a finite set of credence functions $c_1, ..., c_n$.

Background knowledge K: As already explained, orthodox Bayesianism always implicitly conditionalises on a set of background knowledge, K. Even if a credence in a proposition "looks like it is not conditional on anything", like $c(P)$, there is always an implicit conditionalisation on a set of background knowledge K taking place, so $c(P)$ is a shorthand for $c(P \mid K)$. Orthodox Bayesianism usually omits making K explicit, because it is clear that it is always being implicitly conditionalised on (cf. Earman, 1992).

One motivation for fragmented accounts of doxastic attitudes often presented in the literature is limited information access (cf. Elga and Rayo, 2021, Elga and Rayo, 2022, Fleisher, 2023). Agents often have some information available in some circumstances, but different information available in others, and this phenomenon is one of the main motivators for existing accounts of fragmentation, where exactly this phenomenon can be modelled.

For orthodox Bayesianism, the background knowledge K is global; whenever the agent reasons, they use and conditionalise on the same global set of background information. Going to a fragmented version of Bayesian-

ism, the set of background knowledge, K, does not have to, but can be made fragment-relative. This means that fragmented Bayesianism provides the option of modelling limited and fragmented information access. If one chooses to do so, then each fragment / credence function c_n will come with its own fragment-relativised set of background knowledge K_n. Then, contrary to orthodox Bayesianism, there is no global set of background information, but multiple (potentially incompatible) sets of background information, $K_1, ..., K_n$, that are, each respectively, associated with one of the credence functions, $c_1, ..., c_n$, making up the agent's doxastic state.

Note, however, fragmentation does not necessarily require K to be c-relative (fragment-relative). K could still remain a global feature of the agent's doxastic state. Then, the resulting version of fragmented Bayesianism would be closer to orthodox Bayesianism. But fragmentation provides the option of relativising K to fragments. In this case, the resulting version of fragmented Bayesianism would be more permissive, and thereby further away from orthodox Bayesianism.

Having the option to choose whether to keep K global or whether to make it fragment-relative is attractive because it allows one to use and adapt fragmented Bayesianism to model a variety of scenarios. There might be cases where fragmentation is warranted, but not because background knowledge is access-limited. But at the same time, there will be cases where fragmentation is warranted, because the agent has limited information access. And furthermore, rendering background knowledge fragment-relative can help solve central probes of orthodox Bayesianism, such as the problem of omniscience (Elga and Rayo, 2022), and the problem of old evidence (Fleisher, 2023).

(Sub-)Algebra $\mathcal{A}$: Now, one can even go one step further than relativising background knowledge K to fragments. Instead of defining all fragments $c_1, ..., c_n$ on the same algebra, $\mathcal{A}$, one can further depart fragmented Bayesianism from orthodox Bayesianism by allowing for each fragment / credence function c_n to be defined on a different algebra $\mathcal{A}_n$, or different sub-algebras, $\mathcal{A}'_n$, of the same overarching algebra $\mathcal{A}$.

This would further develop the idea presented above of limited information access, where now, it is not just background information that is limited and fragment-relative but information more generally. Committing to relativising the algebra $\mathcal{A}$ to fragments could mean that, eg., $c(P)$ could be defined for one fragment / credence function, but not for another fragment / credence function.

There is another potential application for fragmenting the algebra. A version of fragmented Bayesianism where each credence function comes with a different sub-algebra $\mathcal{A}'$ of the overall algebra $\mathcal{A}$ (or different partitions on W which induce the different sub-algebras), can be used to model levels of description, awareness, or even ontology (List and Pivato, 2015, List, 2017, similar approach also in Leitgeb, 2017). Let me briefly explain how: By allowing for different fragments to correspond to different sub-algebras of $\mathcal{A}$ (or different partitions on W), one can model the fact that an agent is aware of different aspects of the overall algebra in different circumstances[8]. But one can also model more things; List and Pivato (2015) show how objective chance on a high level of description can be compatible with determinism on lower levels of description. For example when tossing a coin, on a high level of description, one seems to be able to ascribe objective chances to the potential outcomes of the coin toss, whereas on a low level "the initial state of the coin determines the outcome" (List and Pivato, 2015). Subsuming both, objective chance, as well as determinism "under one hat", so to speak, has been a challenge, and commonly, it has been argued that objective chances on higher levels of descriptions are actually agent's subjective and uncertain degrees of belief about the events in question. The authors show how both objective chance and determinism are actually compatible, roughly, by showing that "different levels of a system correspond to different specifications of the systems state space [...], at different levels of coarse-graining, which induce different [algebras of propositions] on which probabilities are defined".

A fragmented account of Bayesianism where the algebra $\mathcal{A}$ is fragmented

[8]This is why this could connect fragmentation to the awareness literature, cf. eg. Belardinelli and Schipper (2024).

into different sub-algebras $\mathcal{A}'$ would be the perfect epistemic theory to sit in the back of an account like this, since it models an agent who does exactly what List and Pivato (2015) propose. The agent has different fragments that come with different levels of grain, and thus, different sub-algebras, which makes it possible for the agent to believe that objective chance and determinism can co-exist at different levels of description, which, behind this backdrop, would correspond to / be realised within different fragments of the agent's doxastic state. This would be one attractive application of the fragmented Bayesianism framework where the algebra parameter is relativised to fragments.

Again, by relativising $\mathcal{A}$ to fragments, one makes the resulting version of fragmented Bayesianism depart further from orthodox Bayesianism than its version when keeping $\mathcal{A}$ global / unfragmented.

Notion of justification / confirmation measure M(& threshold t): Another option for making fragmented Bayesianism more or less fragmented is deciding whether to allow for only one global or many, at most one per fragment, confirmation measures / notions of justification, M.

As already explained, orthodox Bayesianism allows one to formulate multiple different and sometimes incompatible measures of justification / confirmation. Another way to increase the level of fragmentation in fragmented Bayesianism is to allow each fragment to have its own unique measure of confirmation / concept of justification. As not all confirmation measures are compatible, this could result in different notions of confirmation / justification being at work for one and the same agent and thus for potentially jointly incompatible propositions to be confirmed / justified in different fragments of the agent's doxastic state. Again, depending on whether one stipulates a specific confirmation measure of the global doxastic state of the agent or whether one allows for a variety of fragment-relative confirmation measures, the resulting account will be more close or more far away to orthodox Bayesianism.

Should one choose a confirmation measure including a threshold for justified credence (or justified belief, if one employs a Lockean account, cf.

Leitgeb, 2017), t, then this will be a further parameter that can either be left globally defined, or fragment-relativised. One can decide whether to globally set a threshold t for all fragments employing a measure of confirmation including a threshold, or whether to allow for thresholds to be set individually in each individual fragment.

Ur-prior U: For orthodox Bayesianism, the ur-prior is the hypothetical probability function that describes the doxastic state of the agent before learning "anything at all" (cf. Fleisher, 2023). Then, the the agent's credences at a certain time are given by conditionalising on the ur-prior, and on all evidence and background knowledge the agent possesses at that time.

Ur-priors also "encode the most basic epistemic standards to which the subject is committed" (cf. Fleisher, 2023). Those are the standards the agent has about how one should properly respond to evidence (cf. Meacham, 2016). Orthodox Bayesians disagree over what set of epistemic standards an agent ought to have. Objective Bayesians argue that an agent should only have a single set of epistemic standards, which means that there is only one way of properly constructing the ur-prior (cf. J. Williamson, 2010). Subjective Bayesianism on the other hand allows for more variation and does not require for there to be only one acceptable set of epistemic standards / one acceptable ur-prior. It is possible to make an inductive risk argument against objective Bayesianism, which will be part of section 5 of this chapter.

For now, the purpose is merely to show that there is multiple ways in which fragmented Bayesianism can take on board commitments about ur-priors. Depending on whether one requires one global ur-prior that constrains all fragments, or whether one allows for some / all fragments to have their own ur-priors, the resulting account of fragmented Bayesianism will be more or less far apart from orthodox Bayesianism. Note that allowing different ur-priors in different fragments goes beyond the position of subjective Bayesianism within orthodox Bayesianism. The reason for this is that subjective Bayesianism says that there can be multiple correct ur-priors, but fragmented Bayesianism with fragment-relativised ur-priors says

that it can be correct to employ multiple different ur-priors at the same time within one and the same doxastic state, which is a much stronger commitment.

Even when allowing for ur-priors to be fragment-relative, however, one can restrict the permissiveness of the resulting account to a certain degree. One can, for example, require "subjective similarity" (cf. Fleisher, 2023). According to this requirement, different fragments are allowed to come with different ur-priors, but they have to be "close enough" to the others. This requirement can be set in place to avoid ad-hoc choices of epistemic standards, and to "ensure that evaluations of evidential responses result from appeal to the subject's general epistemic commitments" (Fleisher, 2023). If one wants the fragments to be "sufficiently similar", then one has to define what "sufficiently similar" means. One way of doing so is by appeal to distance measures from the literature on scoring rules (cf. Joyce, 1998, Pettigrew, 2016, Dunn, 2019, Fleisher, 2023). With the necessity to define a precise measure of what "close enough" means, again, there is the option to define one global measure for what it means, or to define different measures for different subsets of the total set of fragments.

In this section, I have presented a very general version of fragmented Bayesianism and have provided a number of parameters, background knowledge K, (sub-)algebra $\mathcal{A}$, confirmation measure M and threshold t, as well as ur-prior U, that can be either relativised to fragments which results in a more fragmented version of fragmented Bayesianism, or kept global / unfragmented.

The more parameters one chooses to relativise to fragments, the more fragmented the resulting account will be. Furthermore, although it is not the focus of this chapter, the more parameters become fragmented, the more difficult the process of de-fragmenting will become, that is, the process of merging the individual fragments / credence functions into one overall credence function that describes the aggregate doxastic state of the agent.

The account presented here differs from existing accounts of fragmented

Bayesianism (such as Elga and Rayo, 2021, Elga and Rayo, 2022, Fleisher, 2023) in that it is much more general and, unlike them, provides an overview over all the options available within the Bayesian machinery for fragmentation.

5.4 An argument for fragmentation

One might ask – fragmented Bayesianism is a nice formal gimmick that allows for different parameters of orthodox Bayesianism to be fragmented, and thereby weakens the orthodox account. But why does fragmentation make sense? Why would one need a fragmented account of Bayesianism in addition to or instead of the orthodox account?

In this section, I will provide what I call an inductive risk argument for fragmentation. This argument is similar to other arguments for fragmentation already presented throughout this dissertation. The core idea of this argument is that what justifies fragmentation is uncertainty and the inductive risk that is associated with reasoning under uncertainty. In other words: If one wants to provide a model of the doxastic state of an agent who has to reason under (severe) uncertainty, an account like fragmented Bayesianism may be more appropriate than its more orthodox counterpart.

The reason for this is that if the uncertainty is severe enough, then the choice of parameters such as background information K, algebra $\mathcal{A}$, confirmation measure M, or ur-prior U will be epistemically underdetermined. I will now provide an inductive risk argument in favour of fragmented Bayesianism in its most general form, as well as for the fragmentation of each of the parameters presented above, individually.

5.4.1 Argument for fragmentation

The main justification for fragmentation is uncertainty. This argument has already been extensively developed for the qualitative account of justification, and can be transferred to Bayesianism without much modification. Therefore, I will merely provide a brief synthesis here. What will change

is that there are many more options than the qualitative framework for modelling the reasons for why conflicting but equally justified information can occur in the first place.

Let's recall and slightly modify the example from the previous chapter. One the basis of one climate model, a climate scientist forms credence $c(P) = 0.7$ in proposition P, but on the basis of a different climate model, she forms credence $c(P) = 0.3$. The reason for why she results in having conflicting opinions about P is not that either credence is wrong. They were generated on the basis of two highly reliable climate models. The reason is that both models operate on the basis of different sets of background assumptions. And the reason why it is epistemically justified to operate on different sets of background assumptions is uncertainty. The choice of modelling assumptions and background information to consider is underdetermined by the evidence, and thus, these are choices that have to be made under inductive risk, i.e. the risk of being wrong when making such a choice under uncertainty.

The credences the scientist has are a problem for orthodox Bayesianism, because they do not obey the rules of probabilism. Orthodox Bayesianism, and more generally, orthodox epistemology, provides options for dealing with such situations. My argument is that these traditional options are not satisfactory for cases like the climate science case, and that fragmentation is an option for dealing with such a situation that much better respects the epistemic subtleties of such situations.

The first option is suspension of judgment – if you can't decide on what credence to assign proposition P, remain neutral until you can. For Bayesians, this usually means, if you're unsure, randomise until you do not have to randomise. But randomisation, here, would not reflect the fact that these two conflicting credences are based on reliable climate models and that they are well justified.

The second option would be to require one to retain coherence, to commit to either one or the other credence in P. For this strategy, the begging question would be how to choose which credence to throw into the epistemic trash and which one to keep, given that, again, deciding which credence

is epistemically better or worse is underdetermined by the available evidence. Therefore, just ditching either credence would be a completely ad hoc decision, and one would achieve coherence by giving up potentially very valuable information.

Fragmentation is an alternative strategy to those just presented. It does not force one to get rid of either doxastic attitude, while getting rid of the outright incoherence that both orthodox options want to tackle. How? By allowing for the scientist's doxastic state to be made up of two different credence functions, one, c_1 that takes as background information the data generated by one model, and the other one, c_2 that takes as background information the data generated by the other model. The two credence functions can disagree over what credence to assign to a proposition such as P. But each credence function, individually, needs to remain coherent, and adhere to the tenets of orthodox Bayesianism. Thereby, one manages to avoid having an incoherent credence function, but instead "contains" the incompatible information within different credence functions / different fragments.

Fragmentation is an attractive way to deal with the epistemically risky but albeit well-justified information-state the climate scientist finds herself in here.

5.4.2 Arguments for fragmentation of parameters

As presented in section 2, fragmented Bayesianism allows for much fragmentation beyond merely allowing for an agent's doxastic state to consist of multiple credence functions instead of just one.

Background knowledge K: The argument for fragmented background knowledge is only indirectly an argument from uncertainty and more directly an argument stemming from the psychological reality of agents, even epistemically highly able agents – their lack of omniscience and the lack of information or selective availability of information they are often faced with. Part of the reason for this can be uncertainty, but even when dealing

with a subject matter leaving little to no room for uncertainty (such as logic), fragmentation of background knowledge might, in some cases, still be warranted, because it better describes the informational situation the agent finds themselves in.

Fragmenting background knowledge is a way of representing the lack of omniscience agents often have to deal with. It can be a way of representing limited and circumstance-specific access to information which, in the literature on fragmentation, is one of the main motivating factors for employing a fragmented epistemology.

(Sub-)Algebra $\mathcal{A}$: Taking the reasoning on background knowledge one step further, one can relativise the algebra $\mathcal{A}$ to fragments. This would result in there not being one global algebra over which all credence functions / fragments are defined, but that each individual credence function may come with their own individual algebra or sub-algebra. The reasoning for making the algebra fragment-relative could be, beyond the arguments already presented for fragmented background knowledge, that the credence in a certain proposition P should only be defined for the fragment where P is actually epistemically accessible for the agent. If P is not part of the algebra of a fragment, then $c(P)$ should not be defined in that fragment. But, as List and Pivato (2015) and List (2017) show, fragmenting the algebra to have each fragment correspond to a sub-algebra $\mathcal{A}'$ of $\mathcal{A}$ (a specific partition on W) can help model phenomena such as different levels of attention and awareness the agent might (rationally) have with respect to different subject matters or in different circumstances. But it can also model how an agent can entertain different levels of description (fine- or coarse-grained-ness) about an event or phenomenon that seem incompatible with each other, such as the objective chance-description and the determinism description with respect to the event of throwing a dice, explained earlier.

Notion of justification / confirmation measure M(& threshold t): The argument for fragmentation of confirmation measures M is a classic

inductive risk argument. Which confirmation measure M is the right one? This is a much debated questions within orthodox Bayesianism. The reason why this is such an important question is that depending on one's choice of confirmation measure, the resulting versions of Bayesianism look very different, and central arguments of Bayesianism hold or do not hold (eg. the so-called Popper-Miller argument, cf. Fitelson, 1999).

The problem, here, is that choosing a confirmation measure is a choice one has to make under inductive risk. No amount of evidence (of whatever form) can conclusively help one decide which confirmation measure is the right one. There is a gap between the evidence and the choice of a confirmation measure. This means that, purely epistemically, there is no way to decide between the multiple different confirmation measures / notions of justification there are out there.

In his seminal argument, Rudner (1953) argues that inductive gaps like this ought to be filled by using non-epistemic value judgments about the severity of an error when making a (potential wrong) choice. Recall that Jeffrey (1956) has the opposing view that one should just not make a choice, and thereby there is no need to bridge the gap. This strategy would be very unsatisfactory here, because Bayesianism is supposed to be a theory of confirmation / justification. Without choosing any confirmation measure, it would not be such a theory. Hence, this option is not really an option here.

Fragmentation provides another option of looking at this problem. It has been recognised that some confirmation measures are more useful for some purposes, while others are for other purposes (Fitelson, 1999). This suggests a context-sensitivity of confirmation measure, where, depending on the specific epistemic and non-epistemic circumstances (of an agent or situation), one or another confirmation measure might be better or worse. Fragmented Bayesianism turns out to be a perfect formal tool to model this way of dealing with the problem of inductive risk surrounding the choice of confirmation measures. The reason for this is that it allows an agent to work with and accept multiple different confirmation measures at the same time, within different fragments. Thereby, fragmentation manages to

provide a mid-way solution between the classic solutions to the problem of induction, provided by Rudner (1953) and Jeffrey (1956). Fragmentation of confirmation measures is a Jeffreyan solution in the sense that it takes away the need to decide on one single confirmation measure (as Rudner would suggest, using non-epistemic value-judgments). By allowing for fragmentation, multiple different confirmation measures can be permissible at the same time, which is Jeffreyan in nature, because it takes away the need for a definite choice. But it is Rudnerian at the same time in the sense that what justifies employing different confirmation measures simultaneously will be a combination of epistemic and non-epistemic factors making up the circumstances giving rise to the different fragments[9].

Ur-prior U: Similarly, the argument for fragmentation of ur-priors U is a classic inductive risk argument. Recall that the ur-prior is thought to encode the most basic epistemic standards a subject possesses. It is the basic probability function an agent possess before they learn anything at all.

Again, inductive risk tells us that there is no unique set of epistemic standards that is universally better than any other set of standards (this is also what subjective orthodox Bayesians believe). And it can also be seen with the help of the climate science example: Both climate models are reliable and trustworthy. However, one possible reason for them resulting in vastly different predictions about future climate change is that they might encode different epistemic standards which are not necessarily better or worse than each other. Fragmented Bayesianism provides a formal tool to model an agent who (rationally) uses different sets of epistemic standards within different fragments and circumstances.

And note: This whole discussion on inductive risk is not to say, for example, that there is no way to decide between a good and a bad set of epistemic standards. There is – one can distinguish between better and worse sets of standards. Rather, the argument is that, even when

[9]For further examples of such factors, the reader may refer back to the previous chapters.

completely exhausting the epistemic realm, i.e. using all the evidence there is, there will still be competing sets of epistemic standards where no purely epistemic way of deciding as to which one is better or worse is available.

5.5 Comparison to qualitative account

The main difference between the qualitative account of fragmented justification developed earlier on, and the account of fragmented Bayesianism developed here is that the qualitative account provides much less variability regarding the fragmentation of parameters beyond the mere introduction of fragments. Fragmented Bayesianism has a list of parameters that can be chosen to be made fragmented, or not. Thereby, fragmented Bayesianism provides a scale along which the account can be made more or less fragmented, and thereby more or less close to orthodox Bayesianism. There is no such scale for the qualitative account, and there are much less parameters, if any, that can be made fragment-relative beyond the mere introduction of fragments into the account.

Another important difference is that the qualitative account of fragmented justification I developed presupposes a specific notion of justification / confirmation – absolute justification, the most common notion of justification in the literature on qualitative accounts of justification. Instead, the account of fragmented Bayesianism presented here is very general and allows for / is compatible with many different notions of justification / measures of confirmation.

5.6 Conclusion

This chapter has presented the variety of notions of justification that orthodox Bayesianism provides, and developed a very general account – a skeleton– of fragmented Bayesianism. Depending on which parameters of Bayesianism one chooses to relativise to fragment, this skeleton, will become more concrete, and specific to a certain modelling situation, and fur-

thermore, more or less far away from orthodox Bayesianism. The chapter, moreover, provided an arguments as justification for fragmentation, and lastly, provided a brief comparison to the qualitative account of justification developed earlier.

Interesting avenues for further work would be, for example, to investigate in more detail how qualitative and quantitative accounts of fragmentation relate to and differ from each other, in general, beyond just looking at the concept justification in a fragmented setting. Furthermore, and this will partially be done in the following chapter, a potentially fruitful avenue for further research would be to study in detail the aggregation options available for accounts of fragmented Bayesianism. The reason why this would be interesting is that the fragmentation of the Bayesian parameters introduced in this chapter makes aggregation even more difficult than it was in the qualitative setting.

Chapter 6

Fragmentation and Imprecision

Imprecise Bayesianism, as opposed to orthodox Bayesianism, aims to model agents who do not have enough information to form precise credences. Under this account, an agent's doxastic state is allowed to be made up of not one but many different credence functions, to model imprecision in the assignment of credences to propositions. Fragmentation, on the other hand, aims to represent agents with fragmented doxastic states, where fragments are individuated via differences in the agent's epistemic and non-epistemic circumstances, such as differences in accessible information, (methodological) background assumptions, or different practical goals, stakes, etc. The formal model underlying this account is very similar to imprecise Bayesianism: The agent's doxastic state is represented via not one but many different credence functions.

The literature streams on fragmentation and imprecision do not comment on the striking similarity between these two different accounts. This chapter is devoted to presenting and comparing their similarities and differences. The upshot will be that the main difference lies in the interpretation of the formalisms, how these two accounts deal with questions of practical rationality, as well as how they deal with new information entering the doxastic state and updating on that new information.

6.1 Introduction

This chapter presents and compares the accounts of imprecise Bayesianism and fragmented Bayesianism. Imprecise Bayesianism, as opposed to orthodox Bayesianism, aims to be able to model agents who have imprecise credences in propositions, for example, because the evidence they have is not conclusive (enough). To do so, imprecise Bayesianism allows an agent's doxastic state to be made up of multiple credence functions, instead of just one. Fragmented Bayesianism, on the other hand, aims to model fragmented doxastic states of agents – this means that their doxastic state is not represented via a single credence function that obeys the norms of orthodox Bayesianism, but that it may contain multiple different credence functions / fragments that may disagree with each other.

These two different accounts, and the associated streams of literature are very closely related. The formalisms underlying imprecise Bayesianism and fragmented Bayesianism are almost identical. However, there is no explicit connection established in the literature between imprecision and fragmentation. This is the goal of this chapter. It presents both imprecision and fragmentation, and works out important similarities and differences between the accounts. The account of imprecise Bayesianism will be able to be classified as a sub-account of a very general account of fragmented Bayesianism, and the main differences between imprecision and fragmentation will lie in, first, the respective interpretations of the formalism, and second, how both accounts deal with the question of practical rationality, the (need for) aggregation of credence functions, and how they demand the process of incorporating new information into one's doxastic state to take place.

The chapter will be structured as follows: Section 2 recalls the central theses of orthodox Bayesianism as a starting point. Section 3 recalls the account of imprecise Bayesianism. And section 4 presents fragmented Bayesianism, both in a very general version, building on what I developed in the previous chapter, as well as in the most prominent version present in the literature, developed by Elga and Rayo (2021), Elga and Rayo (2022). Sec-

tion 4 also provides a brief overview over the probabilistic opinion pooling literature, since that literature provides the tools for fragmented Bayesianism to deal with the problems it might face in light of practical rationality and decision making. Section 5 compares imprecision and fragmentation with each other, with a special focus on a few aspects. The first focus is the differences and similarities in interpretations of the respective frameworks. Then, I discuss how they deal with the question of practical rationality and the need for aggregation of credence functions. Lastly, I show how the accounts differ when it comes to receiving new information and incorporating it into one's doxastic state.

6.2 Orthodox Bayesianism

Let us recall: The core of orthodox Bayesianism thus consists of two theses, probabilism and conditionalisation (the formulation of those theses roughly follows Vallinder, 2018). Let W be a finite (for reasons of simplicity) and non-empty set of possible worlds. Then $\mathcal{A}$ is a powerset algebra on W, it is the set of all subsets or powerset of W, and it comprises all propositions P within W. Then, Bayesianism is the combination of the following two theses:

1. Probabilism: A rational agent's degrees of belief are represented by a credence function c which is a function from $\mathcal{A}$ into the interval $[0, 1]$. This means that it assigns a credence value between zero and one to all propositions in the logical space. The credence function c respects the axioms of probability theory (Kolmogorov axioms):

   1) $c(P) \geq 0$, for all P
   2) $c(W) = 1$
   3) For incompatible (mutually exclusive) propositions P and Q: $c(P \cup Q) = c(P) + c(Q)$

2. Conditionalisation: A rational agent updates her degrees of belief over time by conditionalising her credence function on all the evidence

she has received. If an agent with credence function c_0 learns new evidence E between t_0 and t_1, then her new credence function c_1 is given as $c_1(P) = c_0(P \mid E)$[1].

Using these two theses together make up the core of orthodox Bayesianism. According to that, a rational agent assigns precise credences to propositions, using a credence function that respects the above conditions on the credence function. Furthermore, when the rational agent receives new information, then the updates her doxastic state by way of conditionalisation.

6.3 Imprecise Bayesianism

But what if the agent does not have conclusive evidence to assign precise credences to a proposition, such as P? Orthodox Bayesianism requires her to do so, but her evidence might not be conclusive enough to make her pick a precise credence with "good enough" reason.

As a response to this worry, many authors, such as Levi (1974), Levi (1980), Levi (1985), Walley (1991), Kyburg and Pittarelli (1996), Joyce (2010), and Augustin (2021) have been defending versions of imprecise Bayesianism. The central idea of imprecise views of Bayesianism is that a rational agent's doxastic state can be represented not by just one but many different credence functions.

To give a concrete example, Joyce (2010) presents imprecise Bayesianism as making the following two claims:

1) Imprecise Probabilism: A rational agent's degrees of belief are represented by a credal state C, which is a set of credence functions $c_1, ..., c_n$. Each $c_i \in C$ assigns a real number x within the interval $[0, 1]$, to each proposition P in the algebra $\mathcal{A}$, $c_i(P) = x$. Each $c_i \in C$ respects the Kolmogorov axioms of probability.

2) Imprecise Conditionalisation: A rational agent updates her degrees of belief over time by conditionalising her credence functions on all

[1]Using the Bayes Theorem, this conditional probability can be reformulated which makes it easier to calculate in practice.

the evidence she has received. If an agent with credal state C_0 consisting of credence functions $c_1, ..., c_n$, learns new evidence E between t_0 and t_1, then her new credal state, C_1, consists of the set of individual credence functions $c_1', ..., c_n'$, where each credence function c_i gets updated via conditionalisation by E: $c_i'(P) = c_0(P \mid E)$[2].

This means that each credence function within the agent's overall credal state behaves just like a traditional Bayesian credence function that obeys the axioms of probability and gets updated via conditionalisation. The difference is that imprecise Bayesianism allows for a set of credence functions. Note that the conditionalisation thesis tells us that if an agent receives new information, this information should be incorporated not into one or some, but all of the agent's individual credence functions[3].

There is an array of literature on philosophical motivations for imprecise Bayesianism. Motivations include the ability of imprecise models of probability to solve the Ellsberg paradox (Ellsberg, 1961). They are also argued to be better models of uncertainty than traditional models of probability (Gärdenfors, 1988, Walley, 1991). Other motivations include allowing for incomplete preferences in the expected utility frameworks (Kaplan, 1983, Levi, 1986, Seidenfeld, Schervish, and Kadane, 1995), or adding to the literature on realism (Arló-Costa and Helzner, 2010)[4].

Furthermore, imprecise probabilities prove a useful tool in many different areas, such as engineering, AI, and climate science (cf. Schoenfield, 2020, p. 969). In "engineering design decisions, it is valuable to explicitly represent the imprecision in the available characterisation of uncertainty by using im-

[2]Using the Bayes Theorem, this conditional probability can be reformulated which makes it easier to calculate in practice.

[3]Joyce (2010) presents an alternative option for imprecise conditionalisation; which simply eliminates those probabilities in the overall credal state C that are incompatible with the new evidence received. But he rejects this option; "this model seems to be based on a confusion between learning that some event has occurred and learning that its prior objective chance of occurring was 1". Therefore, I will not delve into this option further. For more information on why Joyce judges this option untenable, cf. Joyce (2010).

[4]An more extensive overview over motivations for imprecise probabilities can be found in S. Bradley (2019).

precise probabilities" (Aughenbaugh and Paredis, 2006). With regards to AI, "imprecise probability models are needed in inference problems characterised by scarce, vague and conflicting information. So it is not surprising that AI has a long history of research and interest in imprecise probability models" (Zaffalon and De Cooman, 2005, pp. 1f.). Lastly, for the field of climate science, imprecise probabilities prove important, because "one may be unable to supply the required [precise] subjective probabilities that any 'filling in' of the gap between probability ranges and precise probabilities may prove too arbitrary to be a reasonable guide to decision. Policy makers may quite reasonably refuse to base a policy decision on a flimsy information base, especially when there is a lot at stake" (R. Bradley, Helgeson, and Hill, 2017, p. 505)[5]. Especially the comment on imprecise Bayesianism in climate science will become important when comparing fragmented Bayesianism and imprecise Bayesianism, because it is claimed here, that imprecise probabilities might be a more reasonable guide to decision than precise probabilities if the evidential basis for a precise probability is too weak.

The intended interpretation of the formalism of imprecise Bayesianism is most often taken to be that, whenever all individual credence functions (which Joyce (2010) calls (fictional) "committee members" of the overall committee, i.e. overall credal state C) agree on what credence to assign to a certain proposition, then the agent has a precise credence in that proposition. But whenever the individual credence functions / credal committee members disagree over the credence of a proposition, then the resulting overall credence in that proposition is imprecise, it will be the interval between the lowest and the highest assigned credence by any individual credence functions to that proposition. This interpretation sometimes goes under the name "Pareto interpretation" (Smith, 2025).

[5]Another practical motivation for imprecise Bayesianism has been presented by Kyburg and Teng (1999), Kyburg and Teng (2001), where they run simulations of a betting game comparing the performance of a precise Bayesian method vs. an imprecise Bayesian method. Interestingly, over larger numbers of trials, the imprecise Bayesian outperforms the precise Bayesian.

6.4 Fragmented Bayesianism

I will present two variants of fragmented Bayesianism to help showcase how the account relates to imprecise Bayesianism. The first one is the very general account of fragmented Bayesianism I developed in the previous chapter. The second one will be the most prominent account of fragmented Bayesianism present in the literature so far, by Elga and Rayo (2021), and Elga and Rayo (2022). They will differ in an important aspect that will position Elga and Rayo's account somewhere in the middle between my general account of fragmented Bayesianism on the one hand, and imprecise Bayesianism on the other.

6.4.1 General version of fragmented Bayesianism

First, let me recall the general version I developed and presented in the previous chapter. Here, an agent's doxastic state can be represented by multiple credence functions $c_1, ..., c_n$, instead of just one, as orthodox Bayesian epistemology prescribes. Each such credence function obeys the norms of traditional Bayesianism, i.e. probabilism and conditionalisation. This means nothing other than that the same formal rules hold for the general version of fragmented Bayesianism and the account of imprecise Bayesianism. However, with one important exception / addition; as presented in the previous chapter, fragmented Bayesianism allows for the fragmentation of various parameters, such as background knowledge K, (sub-algebra) $\mathcal{A}$, confirmation measure M, and ur-prior U, beyond the "mere" fragmentation of credence functions, $c_1, ..., c_n$.

1) Fragmented Probabilism: A fragmented agent's degrees of belief are represented by a credal state C, which is a set of credence functions $c_1, ..., c_n$. Each individual credence function c_i, may or may not, come with their own, local, set of background knowledge K, (sub-)algebra $\mathcal{A}$, confirmation measure M, and ur-prior U. If any of these parameters is not fragmented / credence-relative, there will be, eg., one

global set of background knowledge K, (sub-)algebra $\mathcal{A}$, confirmation measure M, or ur-prior U, for the overall doxastic state C of the agent that holds for each $c_1, ..., c_n$. Each $c_i \in C$ assigns a real number x within the interval $[0, 1]$, to each proposition P in the respective (or global) (sub-)algebra $\mathcal{A}_i$, $c_i(P) = x$. Each $c_i \in C$ respects the Kolmogorov axioms of probability.

2) Fragmented Conditionalisation:

A fragmented agent updates her degrees of belief over time by conditionalising her credence functions on new evidence. However, unlike for imprecise Bayesianism, this updating process is not necessarily global, which means that the new evidence an agent learns does not necessarily get incorporated into all her credence functions, $c_1, ..., c_n$. Instead, for a fragmented agent, depending on how exactly the fragmentation manifests, new pieces of evidence can be learned only relative to one or some, instead of all, credence functions. Thus, fragmented Bayesianism carries the possibility of local, instead of global, conditionalisation.

Note that conditionalisation and the dynamics of fragmentation more generally have not been explored much in the literature on fragmentation and would be a potential fruitful avenue for further work, the full exploration of which would go beyond the scope of this dissertation.

As explained in the previous chapter, there are multiple ways of making fragmented Bayesianism more or less fragmented. One can decide whether to relativise background knowledge K, the entire algebra $\mathcal{A}$, confirmation measure M (& threshold t), or the ur-prior U to fragments / individual credence functions, in addition to simply allowing for multiple credence functions within the agent's doxastic state. Note that these options are specific to fragmentation, they are not discussed as options to increase imprecision in the literature on imprecise probability frameworks. For imprecise Bayesianism, these parameters are (hidden) "global" parameters, they do not vary for the individual committee members making up the

overall credal state of the agent, but they all share the same "parameter setup".

Let us turn to the interpretation of the fragmentation formalism. As opposed to imprecise Bayesianism, fragmented Bayesianism does not come with a general recipe for how to interpret the formalism with the goal of defining the overall doxastic state of the agent. Here, it is not as simple as checking whether all credence functions agree or whether there is disagreement. The reason for this is that, the more fragmented the account becomes (by i.e. relativising background knowledge, or the entire algebra to fragments), the more complex and non-trivial it will be to find a suitable aggregation mechanism to determine how the overall doxastic state emerges from the individual fragments. For example, if the algebra is fragment-relative, the credence of a specific proposition P might not even be defined on all credence functions, because P might not be part of the algebra of all fragments.

Aggregation of fragments is a highly non-trivial endeavour, as I elaborated on in previous chapters, that will often not yield sensible results, especially the more fragmented the agent will become. This will also be the case when turning from aggregating plausibility orderings to so-called probabilistic opinion pooling, the aggregation of probability functions. Generally, there is a vast array of options for aggregating multiple credence functions to one single credence function, so many more than the one aggregation rule that imprecise Bayesianism subscribes to – the Pareto rule.

Moreover, very interestingly, the Pareto rule is not being discussed as a potential aggregation method for fragments within the context of fragmentation. The reason for this is that, actually, the "aggregation method" that imprecise Bayesianism uses is not a genuine aggregation method – the overall doxastic state of the agent is not one single credence function, after applying the Pareto rule. Whenever individual credence functions disagree over the credence over a proposition, all of those functions "will be part of" the overall doxastic state of the agent where the credence for that proposition is represented as the interval of individual credences.

In contrast, aggregation mechanisms in the context of fragmentation

have the goal of completely getting rid of fragmentation – which means that the "aggregation" method of imprecise Bayesianism would not be admissible, because, here, one does not truly get rid of the multitude of individual credence functions on the level of the overall doxastic state of the agent[6]. I will delve deeper into the application of probabilistic opinion pooling to fragmented Bayesianism in a bit.

6.4.2 Task-specific fragmented Bayesianism

Now, let me contrast this very general and schematic version of fragmented Bayesianism with a specific account developed in the literature, by Elga and Rayo (2021) and Elga and Rayo (2022). Here, Elga and Rayo introduce a parameter distinct from the ones introduced in the previous chapter that will be relativised to fragments – tasks. A fragment, here, is made up of a credence function, and a task[7], based on a practical choice making scenario (cf. also Fleisher, 2023). No parameter of the ones introduced within my account of fragmented Bayesianism are fragment-relative on their account, or even mentioned as a possibility for (further) fragmentation.

Incorporating the notion of a task, and thereby a concept from practical rationality into the account of fragmented Bayesianism is significant for multiple reasons. First of all, it shows that Elga and Rayo conceive of Bayesianism not merely as an epistemic theory of justification, as I do, but also, and maybe predominantly so, as a decision theory, that is, a theory of practical rationality.

Secondly, introducing a dimension from practical rationality into the formal framework brings about an entirely different perspective on questions surrounding aggregation of fragments and the formation of the agent's overall doxastic state. According to Elga and Rayo, for any practical task there will always be at most one credence function relevant – namely the one

[6]In the opinion pooling literature, this difference is sometimes labelled as "compromise vs. consensus" (Dietrich and List, 2016).

[7]Elga and Rayo introduce so-called access tables, where a particular task / choice making scenario corresponds to a single row in the access table. The whole access table then, corresponds, to a conglomerate of different tasks.

already associated with that task or circumstance. Generally, the aggregation of fragments and the discussion on how to generate the agent's overall doxastic state from the set of fragments is motivated via considerations of practical rationality.

Because agents should be able to act in a rational way, I argue elsewhere, de-fragmentation will become necessary. But this is only the case if for some practical decision, multiple fragments that might even contain conflicting information become relevant. Such a situation requires some kind of de-fragmentation, because rational practical action based on a fragmented and potentially overall inconsistent set of information is endangered.

By setting up their fragmentation account in this very specific way, Elga and Rayo prevent this from happening. They do so by eliminating the need for de-fragmentation altogether. The need for de-fragmentation is motivated by practical rationality concerns. But within Elga and Rayo's account practical rationality is a fragment-relative notion. So, to ensure practical rationality, there is no need for an overall consistent and unfragmented doxastic state, and therefore no need for de-fragmentation. Any practical task will only ever require the information provided by the credence function that is associated with the same fragment that the task is associated with. And each individual credence function must adhere to traditional rationality norms, which means that the risk of having incompatible information guiding practical action does not exist here. So, no de-fragmentation process is needed, since de-fragmentation becomes necessary whenever the overall fragmented state contains information that could lead to practical irrationality.

Very importantly, thereby, the questions of how to define *purely epistemic* notions of belief and justification move to the background – because to define those notions, one would need to consider fragmented Bayesianism as an epistemic theory of justification first, before delving into practical rationality.

This account differs, both from the general account of fragmented Bayesianism I developed in earlier work, as well as from imprecise Bayesianism. It differs from the general account of fragmented Bayesianism in that this ver-

sion of fragmented Bayesianism is a fragmented decision theory, because, instead of relativising any of the epistemic parameters presented above to fragments, it chooses to make practical rationality a fragment-relative notion. It thereby gets rid of the problem of de-fragmentation. However, the question remains how to associate credence functions with specific tasks, and how to make sure that fragments remain distinct from each other, since tasks could potentially overlap or be closely related, which would make the problem of de-fragmentation reappear from the backdoor.

It differs from imprecise Bayesianism, because it does not require the consideration of what constitutes the overall doxastic state of the agent – which is the central concern of imprecise Bayesianism. Imprecise Bayesianism introduces multiple credence functions to be able to better represent the overall doxastic state of an agent when having, for example, inconclusive evidence. Elga and Rayo's fragmented Bayesianism introduces multiple credence functions, each relativised to a task, not to represent the overall doxastic state of the agent, but to truly fragment it and keep it fragmented, even for and especially in light of, practical rationality.

6.5 Aggregation of credence functions

As one can already see, how the different accounts deal with questions surrounding practical rationality and aggregation of credence functions will be one of the main factors them that sets them apart from each other.

As already indicated, there is a stream of literature that deals with how "several individual probabilistic opinions [should] be unified to collective probabilistic opinions" (Dietrich and List, 2016) – probabilistic opinion pooling (cf. eg. Genest and Zidek, 1986), originating from and closely related to the broader literature on social choice and judgment aggregation already talked about in this dissertation. Perhaps surprisingly, there exists literature on the aggregation of imprecise probabilities, too (Nau, 2002, Benavoli and Antonucci, 2009, Stewart and Quintana, 2018, Quintana, 2024). I will skip this topic here, although a comparison would be interesting and

shall be subject to further work, as aggregation of imprecise probabilities does, at least prima facie, conflict with the central philosophical aim of imprecise Bayesianism, which is to allow for imprecision on the overall doxastic state of the agent which could be endangered by applying "genuine" aggregation rules, as we will see shortly.

I will now provide a brief overview over common probabilistic opinion pooling strategies that can be applied to de-fragmenting fragmented Bayesianism. First, probabilistic opinion pooling. Similar to what has been already presented on the aggregation of plausibility pre-orders, here, too, the central message is that there exist many different options for probability aggregation, and each option will satisfy different axioms / desirable properties for the aggregation mechanism, can be justified under different conditions and will serve different purposes. As already argued for in the previous chapter, therefore, non-epistemic value judgments will become necessary if one wants to choose an appropriate de-fragmentation rule for credence functions in a certain circumstance.

Amongst the most common opinion pooling options for probabilities are linear, geometric and multiplicative pooling rules. Linear pooling goes back to Stone (1961), DeGroot (1974), and back to Laplace. For illustrative purposes, here, individual credence functions are mapped onto a collective credence function / overall credal state, C, in the following way, for each proposition P in the algebra:

$$C_{c_1,...,c_n}(P) = w_1 c_1(P) + ... + w_n c_n(P),$$

where $w_1, ..., w_n$ are fixed non-negative weights with sum-total 1. "The class of linear pooling functions includes a variety of functions, ranging from linear averaging with equal weights, where $w_i = \frac{1}{n}$ for all i, to an expert rule or dictatorship, where $w_i = 1$ for one individual and $w_j = 0$ for everyone else [weighted linear pooling]" (Dietrich and List, 2016).

Next, geometric pooling rules: Here, individual credences in a proposition (or world, because geometric pooling rules are defined on worlds rather than propositions, Dietrich and List, 2016), are multiplied instead of using the sum, and then the nth root operation of the product of the individual

credences is applied, where n amounts to the number of credences being averaged. Again, geometric pooling can be weighted or unweighted, similar to linear pooling.

Formally, geometric pooling looks like this: Individual credence functions are mapped onto a collective credence function / overall credal state, C, in the following way, for each world w in the algebra:

$$C_{c_1,...,c_n}(w) = k\,[c_1(w)]^{w_1} * ... * [c_n(w)]^{w_n}, \text{ where}$$

$w_1, ..., w_n$ are fixed non-negative weights with sum-total 1 and k is a normalization factor, given by

$$k = \frac{1}{\sum_{w' \in \mathcal{A}} [c_1(w')]^{w_1} ... [c_n(w')]^{w_n}}$$

This normalization factor ensures that the total sum of the collective probabilities is amounts to one[8].

Lastly, there are multiplicative pooling rules, which are used when some of the individual credence functions are based on private information, i.e. information only accessible to some credence functions / fragments and not to others (eg. when background knowledge K is fragment-relative). Multiplicative pooling rules are very similar to geometric ones in that they also use the product of individual credences and then renormalise them with a renormalization factor, but they do not apply the geometric average, i.e. they do not take the nth root of the product of individual credences.

Formally, multiplicative pooling looks like this: Individual credence functions are mapped onto a collective credence function / overall credal state, C, in the following way, for each world w in the algebra:

$$C_{c_1,...,c_n}(w) = k\,c_0 c_1(w) * ... * c_n(w), \text{ where}$$

c_0 is some fixed probability function, called the calibrating function, and k is a normalization factor, given by

$$k = \frac{1}{\sum_{w' \in \mathcal{A}} c_0(w') c_1(w') ... c_n(w')}$$

[8]Taking w_1,...,w_n as the exponent / superscript is identical to applying the root operation, since, eg. $\sqrt[2]{x_1, ..., x_n} = x_1^{\frac{1}{n}}, ..., x_n^{\frac{1}{n}}$.

Each of these types of pooling approaches provide many different specific tokens of pooling functions, so the overall space of possible pooling functions for credences is very big. Different ones satisfy different axioms and can be justified on different kinds of grounds. There is no one size fits all approach to choosing the appropriate pooling function in a particular scenario, and the choice will depend on the goals one wants to achieve with the pooling function in the respective circumstance.

Lastly, recall how fragmented Bayesianism contains a variety of parameters that can be chosen to be kept "global" like in orthodox Bayesianism, or to render fragmented, such as background knowledge K, the algebra $\mathcal{A}$, the ur-prior U, or confirmation measure / notion of justification M. Perhaps obviously, the more of these parameters will be fragmented, the more difficult and non-trivial the quest for an appropriate de-fragmentation method will become. The reason for this is that, if such parameters are fragmented, the de-fragmentation mechanism will not just have to send multiple individual credence functions in to one global credence function, where the individual credence functions agree over K, $\mathcal{A}$, M or U, as it is common in the opinion pooling literature. Rather, if the individual functions disagree over K, $\mathcal{A}$, M or U, the de-fragmentation mechanism will not just have to aggregate the credence functions themselves, but find ways to merge multiple sets of background knowledge into one, or multiple algebras into one, or multiple confirmation measures into one, etc. How exactly this could look like would go beyond the scope of this chapter and dissertation, but would be an interesting avenue for further work.

6.6 Comparison

I will now delve deeper into the comparison between imprecise Bayesianism and fragmented Bayesianism.

Both accounts share a very similar formalism – agents' doxastic states are being represented by sets of credence functions. For imprecise Bayesianism, the goal is to model an agent with imprecise credences. For fragmented

Bayesianism, the goal is to model a mentally fragmented agent.

But, although at first glance one might think that the formalisms are identical; they are not. For fragmented Bayesianism there is a list of parameters, such as background knowledge K or the algebra $\mathcal{A}$ etc., that can be chosen to be fragment-relative or to be kept global. For imprecise Bayesianism, these parameters are not credence-relative, rather they are kept as global or "unfragmented" parameters.

Another difference in the basic setup of the accounts concerns the conditionalisation thesis. For imprecise Bayesianism, information updating is "global" in the sense that if the agent receives a new piece of information, every individual credence function $c_1, ..., c_n$ will be updated with this new information. For fragmented Bayesianism, however, this is not necessarily the case. Here, a piece of information might be introduced only to one or a few of the total fragments of an agent. So, conditionalisation can be a highly local matter, and may look very different for different fragments, depending on the information they receive and incorporate, respectively.

For imprecise Bayesianism, the overall doxastic state of the agent is constructed via the "Pareto interpretation" – whenever individual credence functions agree over the credence to assign to a proposition P, then the credence of the overall agent in that proposition is precise. But whenever the individual credence functions disagree over what precise credence to assign to a proposition Q, then the credence of the overall agent in that proposition is imprecise – it is given via the range between the lowest assigned credence and highest assigned credence in the pool of credence functions. In the literature on imprecise Bayesianism, it is argued that, whenever the information base is too weak to warrant precise credences, imprecise credences may be a better basis for practical rationality.

For fragmented Bayesianism, there is no general rule on how to construct the overall doxastic state of the agent. For Elga and Rayo's task-dependent version of fragmented Bayesianism, there is no relevant overall doxastic state, because practical rationality is a fragment-relative matter. As I presented throughout this dissertation, aggregation of fragments becomes relevant in light of questions about practical rationality. But acting

based on one's doxastic state is a fragment-relative endeavour for Elga and Rayo, so the need to get rid of fragments to ensure practical rationality is being eliminated on their account.

But for the more general version of Bayesianism presented above and in the previous chapter, as well as for the account of qualitative fragmented justification developed in the first part of this dissertation, matters of aggregating fragments look much more difficult. Here, multiple fragments might become relevant for one single practical decision which is why the aggregation of fragments becomes a relevant topic. In chapters 3 and 4 (albeit in a qualitative framework), I have provided examples that show that there can be cases where fragmentation is rationally warranted, but in light of a practical decision, multiple fragments will be relevant. There is no universally correct way to aggregate fragments, whether they take the form of plausibility pre-orders, or now, credence functions – choosing an aggregation method is a choice that has to be made under inductive risk, and that will depend on how one needs the resulting information to be "packaged", in order to be useful for purposes of decision-making.

I will now delve deeper into some more specific lines of comparison – first, the philosophical interpretations of the two frameworks, and second, the respective options for credence aggregation, as well as implications for practical rationality that follow from imprecise Bayesianism and fragmented Bayesianism. Lastly, I will briefly touch on another dimension – how the respective accounts deal with the dynamics of information change.

6.6.1 Interpretations of the formalisms

Imprecise Bayesianism aims to model an agent who has enough information to assign some kind of credence to a proposition, but not enough information to assign a precise credence to that proposition. The agent's evidence does not warrant a specific precise credence x in a proposition P, but instead it is only strong enough to warrant a credence range $[x, y]$ in that proposition.

Fragmented Bayesianism aims to model an agent who is fragmented.

There is much choice in terms of the sense in which the agent is fragmented; but generally, the way I perceive of fragmentation, the reason can be summarised as – the agent has very good but conflicting evidence which makes it impossible to employ a orthodox Bayesian credence function on pain of explosion.

At first, these two stories sound very different. For imprecision, there is not enough information to be able to apply orthodox Bayesianism, for fragmentation, there is too much / competing information to be able to apply orthodox Bayesianism.

I will now use an example to show how these two stories might not be so different after all. "A detective might discover evidence that lends weight to the view that the perpetrator was female (suppose that a certain parrot was observed to squawk and that it has been noted in the past that the parrot only squawks in the presence of females) but lends no weight to the view that the perpetrator was Apate in particular (likewise Brimo, Circe and all the other possible female perpetrators). [...] The imprecise probabilist models such situations as follows. Suppose we want to say that the evidence lends probability 0.4 to the claim that a woman did it while lending no particular probability to the claim that Apate (Brimo etc.) did it. The model involves a set of probability functions, all assigning 0.4 to the set of all women (in the sample space of possible perpetrators) and between them assigning all possible values (compatible with there being probability functions) to the singletons Apate, Brimo etc." (Smith, 2025, p. 4f.). This is as far as the imprecision story goes. But, one might ask, why, in fact, are we so unsure about who exactly did it while being sure enough to quantify our uncertainty well enough to assign a precise credence with regards to whether a woman did it or not? What are the reasons for this? What does the evidence look like leading us to assign a credence range to the propositions expressing that anyone in particular did it, in comparison to the evidence leading us to assign a particular credence value to the proposition expressing the fact that a woman did it?

On the fragmentation view, one can go one step further and ask, why would the agent be unsure exactly? She has insufficient or potentially

conflicting evidence with regards to who exactly is the perpetrator. It might be that the agent assigns credence 0.7 to the proposition that "Apate is the perpetrator", because of the evidence that the agent has seen her perpetrate many times before. But then, the agent also has evidence leading her to believe that Apate is on holiday out of the country and cannot physically have been the perpetrator. So, she also wants to assign 0.3 to that proposition. The agent is unsure about which credence to assign to the proposition *for certain reasons*. The agent might also have differing practical motivations and stakes associated to her uncertainty in assigning a precise credence value to that proposition. Maybe the agent is Apate's friend and wants to stay her friend and thereby dos not want to falsely accuse her of something she might not have done (risk-averse), and thereby assign a relatively low credence. On the other hand, the agent might know that saying something false in front of the police is against the law, and since she did see Apate be the perpetrator before, she might want to assign a higher credence[9].

What fragmentation essentially does is allowing for different reasons for why the agent might be unsure to be encoded explicitly into the framework. Any proposition, P, will be assigned a credal value only relative to a certain fragment with certain fragment-specific supplementary information (such as background information etc.).

This example intends to show that, actually, in a sense, the imprecision and the fragmentation stories are not that different. In the imprecision framework, putting on "fragmentation-glasses" can help identify what exactly might be the reason for the choice of credence range constituting the imprecision. Why are the lower and upper bound chosen as they are? In this sense, the imprecision framework can be supplemented by fragmenta-

[9]One more consideration. One might think that the question whether the perpetrator is female and the question who exactly is the perpetrator are two different questions (with two different subject matters) whose potential answers are different partitions over the logical space. Because this is the case, one might argue, one cannot evaluate both questions relative to the same credence function. This suggests that the agent has a fragment / credence function that assigns credences to answers to the first question, and a different fragment that assigns credences to answers o the second question.

tion by way of explicitly encoding into the framework the reason for the imprecision.

This shows that, actually, the stories behind the imprecision and the fragmentation framework are not that different. Fragmentation manages to go one step further in asking and encoding the answer to why there is imprecision with respect to a certain proposition.

But in a different sense, the interpretations of the framework are very different. Imprecise Bayesianism aims to allow for imprecision on the overall level of the doxastic state of the agent. It does so by stipulating that, whenever individual credence functions disagree over which exact credence, the overall credence of the agent in that proposition is imprecise. The relevant interpretatory perspective of imprecise Bayesianism is not on the level of individual credence functions, but on the "aggregate" level, the overall doxastic state of the agent, computed via the Pareto interpretation.

For fragmented Bayesianism, it is definitely not the goal to model imprecision in this way. On the overall level of the doxastic state of the agent, there should *not* be imprecision. And, furthermore, this is not the primarily relevant interpretatory perspective; here the primarily relevant perspective is on the level of individual credence function, individual fragments.

The overall aggregate doxastic state of the agent only becomes relevant when they face decisions that involve multiple fragments that potentially stand in conflict with each other. But here, the Pareto rule would not be a suitable "aggregation" mechanism, since here, on the aggregate level, the goal would be to get rid of the disagreement between individual credence functions (which the Pareto rule cannot achieve), whereas imprecise Bayesianism wants to allow for explicitly this.

6.6.2 Implications for practical rationality

I will now compare the different ways in which the imprecision and the fragmentation framework deal with issues surrounding practical rationality.

Let me start with imprecise Bayesianism – recall how within the litera-

ture on imprecise probabilities it is argued that, when the information base is very limited, then imprecise probabilities might be a more reasonable guide to decision than precise probabilities. This is a Jeffreyan argument relating to how to deal with the problem of inductive risk (cf. Jeffrey, 1956). Instead of eg. a scientist potentially making the wrong call when trying to assign a precise credence to a proposition while having an insufficient evidential basis to make that decision, this argument says that is best to not make this choice, but rather communicate a range of credences in which the proper credence most probably lies. This, according to Jeffreyan reasoning, is a much better basis for policy makers to then make practical decisions on than precise credences, because it makes it possible for the scientist to not have to employ non-epistemic value judgments to bridge the evidential gap to make the call to assign a precise credence. It leaves the employment of non-epistemic value judgments to the policy maker, which is good, because it is policy makers and not scientists who should / are allowed to employ non-epistemic value judgments, one might think.

There are two problems with this reasoning, though. First, actually, even when allowing for ranges of credences, instead of precise ones, the scientist / agent has to make decisions under inductive risk, i.e. the risk of being wrong when making the call. Those decisions are, for example, to set an upper and lower bound for the credence range. Intuitively, at least, it is not clear how making this decision is any different from assigning a precise credence in the first place. In any case, there is an inductive step between the evidence and setting an upper and lower bound that needs to be taken. So the Jeffreyan reasoning of managing to avoid these kinds of decisions by being less committal and using imprecise credences does not fully go through.

The second problem is whether imprecise probabilities can be of any help to practical rationality at all. Of course, important decisions should not be made when the informational basis is too weak. But how should a practical decision be made on the basis of imprecise probabilities instead of precise ones? At least prima facie, it seems much more difficult for eg. a policy maker to base their practical decisions on an information state

containing more rather than less uncertainty, even if the advantage of that is for the scientist coming up with the information state to not have to employ (or may have to employ fewer) non-epistemic value judgments to decrease uncertainty.

Before coming to practical rationality in the context of fragmentation, because this will be an important contrast, let me note that within the imprecision framework, the Pareto interpretation that provides the recipe for generating the overall doxastic state of the agent from the individual credence functions actually does not get rid of the existence of a set of credence functions even on the aggregate level of the overall agent. To construct the overall doxastic state of the agent, here, it is not needed to aggregate the set of individual credence functions into a single overall credence function. All individual credence functions are kept in the overall doxastic state of the agent, and wherever they agree on a credence for a proposition P, the overall credence is precise, and whenever they disagree, the overall credence is imprecise and given via a range over precise credences.

In fact, getting rid of the individual credence functions on the overall level of the agent would be contrary to the goal of imprecise Bayesianism, because it would then not allow for there to be multiple credence functions assigning different credences to the same proposition, which is exactly what is needed to be able to have imprecision in the doxastic state. This will be an important difference to fragmentation, since there, the primary goal is not to model imprecision on the overall doxastic state of the agent, but for the overall doxastic state after having undergone a de-fragmentation process to be a reasonable guide to practical rationality.

Now, let me come to practical rationality in the context of fragmentation. I will first discuss Elga and Rayo's version and then my more general version of fragmented Bayesianism. Elga and Rayo's version makes tasks fragment-relative. Thereby practical rationality is not a question that is to be asked with respect to the overall (aggregate) doxastic state of the agent, but a fragment-relative notion. Having multiple different credence functions that potentially conflict with each other is, therefore, not a problem for practical rationality, because whether an action is rational can only ever be evaluated

relative to one single fragment, i.e. one single credence function that obeys to the norms of orthodox Bayesianism. Hence, practical rationality will be no more difficult to achieve here than within orthodox Bayesianism.

But for the more general version of fragmented Bayesianism I developed, practical rationality is a global notion, at least generally. One can imagine many scenarios where an agent is faced with a decision that depends on potentially conflicting information that is contained in multiple different fragments. Here, on pain of practical irrationality, aggregation of fragments becomes necessary. Elsewhere, I have already argued for the necessary employment of non-epistemic value judgments to decide on a suitable aggregation mechanism, since aggregation of credence functions is a highly non-trivial matter and involves a choice that has to be made under inductive risk. Here, in many cases, such as the climate science example discussed elsewhere, de-fragmentation is demanded by practical rationality. This sharply distinguishes fragmentation from imprecision – in the former unlike the latter, genuine de-fragmentation, i.e. genuine aggregation of credence functions will be necessary, at least in some cases.

6.6.3 Dealing with new information

Lastly, departing from this discussion on practical rationality, let me say a few words on how both accounts differ when it comes to what happens when the agent receives new information (i.e. in a dynamic setting where information changes and has to be updated).

Recall how the conditionalisation process happens for imprecise Bayesianism when an agent receives a new piece of information. This new piece of information gets incorporated into each individual credence function, $c_1, ..., c_n$, and each credence function, individually, ought to follow the conditionalisation process to incorporate this new new piece of information. Conditionalisation here, happens locally, within each credence function individually. However, it does not happen truly locally, since "accessing" the new piece of information to be incorporated into each individual credence function is thought to be a global process. New pieces of information are

assumed to be accessible to *all* individual credence functions in the same way. It is not the case here, that there might be a new piece of information that is accessible to some individual credence function c_i, but not to another one, c_j.

Conditionalisation can look different from this for fragmented Bayesianism. The way I constructed the account of fragmented justification in the previous chapter, the account is a skeleton that allows for multiple ways in which the conditionalisation process in response to an agent receiving new information can take place.

The process of conditionalisation can look similar to how it looks for imprecise Bayesianism, a new piece of information being incorporated in the same way into each and every fragment / individual credence function. But this must not necessarily be the case. If information is not a global notion, but rather, as is one of the main motivations for fragmentation accounts, restricted or selectively accessible, then conditionalisation will look different, here. If a new piece of information is only accessible for some, but not other fragments, only those fragments / credence functions that the new piece of information is accessible to would be able to include the new piece of information by conditionalising on it.

Elga and Rayo's task-relative version of fragmented Bayesianism is an extreme case of this. Different fragments correspond to / are active in different circumstances. If an agent learns a new piece of information in a specific circumstance, then that piece of information can only be conditionalised on in the specific fragment related to that circumstance, but not in any other.

This means that, depending on the specific way fragmented Bayesianism is fragmented ("Is information access restricted?", etc.), the process of conditionalising, and more generally dealing with new information will look differently. It would be an interesting avenue for further research to delve deeper into dynamic aspects of fragmentation in general, and (both qualitative and quantitative) fragmented justification in specific, to formulate norms for the dynamics of information change within fragmented models of belief and justification (both qualitative and quantitative).

Let me make a last point with respect to the dynamics of information change that distinguishes imprecise from fragmented Bayesianism. For the imprecise probability framework, receiving new information has the potential to decrease imprecision. If the agent receives new information, this has the potential to better bridge the evidential gap one has to take to decide on a precise credence. Relevant new information, here, may either decrease the range of imprecision regarding a specific proposition, or eliminate imprecision entirely.

For fragmentation, things are looking differently. Here, new information does not guarantee any kind of decrease in the severity of fragmentation the agent's mental state exhibits. If information itself is a fragment-relative notion, then the agent only ever receives new information relative to a fragment, which cannot affect the overall level of fragmentation. If information is a global notion, then, still, new information does not in any way automatically trigger de-fragmentation of any kind.

6.7 Conclusion

One may conclude by saying that, considering the formal apparatus, imprecision can be viewed as a specific variety of fragmentation, where the "aggregation" to the overall doxastic state happens via the Pareto rule. However, semantically, there are significant differences; for imprecision the relevant level of description / perspective is the overall doxastic state of the agent, and for fragmented Bayesianism, the relevant level of description / perspective is the individual fragments.

Furthermore, the "Pareto rule" is not a genuine aggregation rule, unlike within fragmentation, where the goal of de-fragmentation is to actually get rid of individual credence functions disagreeing over certain matters.

At the same time, although one might think both frameworks model very different phenomena, they are much closer related than one might think. Fragmentation can, in some sense, be viewed as taking imprecision one step further and asking and incorporating into the framework the reason(s) for

the imprecision.

I have shown how the formal frameworks, albeit being very similar, differ from each other. I have also shown how imprecise Bayesianism and fragmented Bayesianism deal with practical rationality in different ways and that the Jeffreyan argument in favour of preferring imprecision over precise probabilities whenever the information base is limited, is problematic.

There is much possible further work on the topic of relating the literature on imprecision and fragmentation. One example is the topic I only briefly touched upon in the end – to compare in more detail how both accounts deal with dynamic information change. Furthermore, it would be interesting to investigate further how concrete aggregation rules could look like for fragmented Bayesianism, especially when further parameters such as background knowledge K, (sub-)algebra $\mathcal{A}$, confirmation measure M or ur-prior U are fragmented, as well, and how such rules would differ from aggregation rules being recently developed for imprecise probabilities, such as Stewart and Quintana (2018) and Quintana (2024).

Chapter 7

Conclusion

This dissertation, first, developed and applied a novel qualitative account of fragmented justification. Within this account, justification is modelled via fragment-relative non-monotonic justification relations between evidence and propositions in need of support.

The first application of the account was a novel solution to the preface paradox that improves on existing (fragmentation) solutions to the paradox. The second application was the IPCC assessment report on climate change, where scientific experts about the climate seem to be justifiedly believing in an overall inconsistent set of propositions. Fragmented justification could usefully be applied to these situations, because the toolkit from orthodox epistemology fails to optimally capture and honour the epistemic intricacies of situations like these. I explained how fragmented justification is rational in these cases, and I developed a set of precautionary norms of fragmentation.

Furthermore, I talked about how sometimes, even if fragmentation is epistemically rational, de-fragmentation will become necessary nonetheless in light of practical rationality. I employed tools from social choice and judgment aggregation to the case of de-fragmentation within a single agent and showed how the result from social choice that there will be no universally applicable / desirable de-fragmentation / aggregation rule also applies to the fragmentation scenario.

In chapter 3, concerning the preface case, I showed how de-fragmentation is very difficult, but ultimately not necessary on the grounds of practical rationality. However, in chapter 4, I showed that de-fragmentation will be necessary and argued that due to the inductive risk associated with choosing a de-fragmentation rule, non-epistemic value judgments will have to play a role when choosing a suitable de-fragmentation rule in a specific circumstance.

In chapters 5 and 6, the dissertation turned to quantitative approaches of justification, and specifically the Bayesian approach to justification. Here, I developed a novel and very general version of fragmented Bayesianism, provided an inductive risk argument for the fragmentation of Bayesianism and its parameters, and I compared the framework of fragmented Bayesianism with that of imprecise probabilities.

This dissertation is an addition to the literature on justification, by providing the first explicitly fragmented account of fragmentation. I showed how existing orthodox accounts of justification are all implicitly based on a non-fragmented picture of belief and rationality. The dissertation adds to the literature on fragmentation, by providing an account of fragmented justification to the already existing accounts of fragmented belief and knowledge.

It also draws connections between justification and belief in relation to fragmentation, and helps to answer some open questions in the fragmentation literature, most prominently the question of whether and why fragmentation can be rational. This dissertation shows how fragmentation is not just a "necessary evil" as it is often portrayed in the fragmentation literature, an account that is more descriptively adequate but fails to be normatively adequate. It shows that fragmentation, sometimes, is the rationally superior thing to do. This is significant, because it repositions fragmentation as a properly normative account.

In addition, as already mentioned, it provides a new and attractive solution to the preface paradox, and thereby contributes to the already vast literature of solutions to this paradox.

Furthermore, it provides a new application for the tools from social

choice judgment aggregation and probabilistic opinion pooling. These tools were developed and are usually applied to cases where the goal is to form a collective opinion from a group of individual agents. Fragmentation, and especially de-fragmentation, is a new application of this machinery, since here, the aggregation is not from individual opinions to a group opinion, but instead, from fragments within a single agent to the overall unfragmented mental state of that single agent.

It adds, moreover, to the literature on Bayesianism, by providing an account of fragmented Bayesianism, with a special focus on the parameters within orthodox Bayesianism, such as background knowledge, the algebra, the confirmation measure, and ur-prior, that can be chosen, each, to be fragmented or global, which will render the overall version of fragmented Bayesianism more or less intensely fragmented, and more or less close to orthodox Bayesianism. I explained how, the more parameters will be fragmented, the more difficult it will be to formulate appropriate de-fragmentation rules for the respective versions of fragmented Bayesianism. I also provided an inductive risk argument for fragmented Bayesianism, with arguments for and example cases of the fragmentation of the individual parameters.

Lastly, the dissertation adds to the literature on imprecise probabilities, by connecting the literature on imprecise probabilities to fragmentation. The reason why this is fruitful and interesting is that the respective formal frameworks are strikingly similar to each other, btu also exhibit important differences from each other.

As it is the beauty with philosophy, and research in general, there are many open questions and potential avenues for further research on the topic of this dissertation. Let me give an overview over some of them.

First, it would be interesting to further investigate the dynamics / diachronic aspect of fragmentation generally, and fragmented justification specifically. How should fragments react to new information? How should the updating process take place? This would be interesting to investigate both for the qualitative, as well as for the quantitative account. In this context, it would also be interesting to see how the answers to these ques-

tion relate to de-fragmentation. Can receiving new information facilitate a (potentially partial) process of de-fragmentation, or are the dynamics of information change and the question of de-fragmentation entirely distinct from each other?

Related but separate, it would be interesting to further develop the connection between imprecision and fragmentation, especially with respect to the dynamics of information change, but also regarding the aggregation of credence functions.

Concerning aggregation of fragments or de-fragmentation, one might want to develop a more sophisticated and fuller account of de-fragmentation, within the quantitative scenario with special focus on the de-fragmentation of the Bayesian parameters, potentially with some formal results distinct form the judgment aggregation and social choice literature.

Logicians might also find it interesting to develop logics of fragmented justification and or belief.

Another interesting avenue for further research would be to look at fragmented agents in social settings. What if multiple fragmented agents engage in forming a collective opinion? Answering this question will require, both, orthodox judgment aggregation / opinion pooling for merging the opinion of many individual agents to one collective opinion, but with the fragmentation of each individual agent, there is a second layer that will complicate the matter. It will turn into a problem about the aggregation of sets of fragments (each set representing one agent) into a collective state that could be either fragmented itself or unfragmented.

Furthermore, and lastly, it could be interesting to investigate whether it might be possible to reframe fragmentation accounts as ways of providing so-called bridge principles between the principles of logic and / or formal epistemology and traditional epistemology[1]

Of course, this list is not exhaustive, and I am looking forward to feedback and further ideas on how the ideas and topics developed in this dissertation could be further expanded, and what other open questions that I

[1]For a recent example showcasing the debate cf. eg. Field and Jacinto (2022).

might have not considered might arise in the future.

Bibliography

Alchourrón, C., P. Gärdenfors, and D. Makinson (1985). "On the logic of theory change: Partial meet contraction and revision functions". In: *The Journal of Symbolic Logic* 50.2, pp. 510–530.

Anderson, C. and K. Kellam (1992). "Belief Perseverance, Biased Assimilation, and Covariation Detection: The Effects of Hypothetical Social Theories and New Data". In: *Personality and Social Psychology Bulletin* 18, pp. 555–565.

Annis, D. (1978). "A Contextualist Theory of Epistemic Justification". In: *American Philosophical Quarterly* 15.3, pp. 213–219.

Arló-Costa, H. and J. Helzner (2010). "Ambiguity Aversion: The Explanatory Power of Indeterminate Probabilities". In: *Synthese* 172.1, pp. 37–55.

Arrow, K. (1950). "A Difficulty in the Concept of Social Welfare". In: *Journal of Political Economy* 58.4, pp. 328–346.

Artemov, S. (2008). "The Logic of Justification". In: *Review of Symbolic Logic* 1.4, pp. 477–513.

— (2019). *Justification Logic: Reasoning with Reasons*. Ed. by Melvin Fitting. Cambridge University Press.

Artemov, S., M. Fitting, and T. Studer (2024). "Justification Logic". In: *The Stanford Encyclopedia of Philosophy*. Ed. by E. N. Zalta and U. Nodelman. Fall 2024. Metaphysics Research Lab, Stanford University.

Atkinson, D. (2012). "Confirmation and Justification. A Commentary on Shogenji's Measure". In: *Synthese* 184.1, pp. 49–61.

Aughenbaugh, J. and C. Paredis (2006). "The Value of Using Imprecise Probabilities in Engineering Design". In: *Journal of Mechanical Design* 128.4, pp. 969–979.

Augustin, T. (2021). "Statistics with Imprecise Probabilities – A Short Survey". In: *Uncertainty in Engineering. Introduction to Methods and Applications*. Ed. by L. Aslett, F. Coolen, and De Book J. Springer, pp. 67–79.

Aumann, R. (1962). "Utility Theory without the Completeness Axiom". In: *Econometrica* 30.3, pp. 445–462.

Baccini, E. et al. (2023). "The Wisdom of the Small Crowd: Myside Bias and Group Discussion". In: *Journal of Artificial Societies and Social Simulation* 25.7, pp. 1–7.

Bacon, A. (2023). "A Theory of Structured Propositions". In: *Philosophical Review* 132.2, pp. 173–238.

Baltag, A., V. Fiutek, and S. Smets (2016). "Beliefs and Evidence in Justification Models". In: *Advances in Modal Logic, Volume 11*. Ed. by L. Beklemishev, S. Demri, and A. Máté. CSLI Publications, pp. 156–176.

Baltag, A., B. Renne, and S. Smets (2012). "The Logic of Justified Belief Change, Soft Evidence and Defeasible Knowledge". In: *Logic, Language, Information and Computation. WoLLIC 2012. Lecture Notes in Computer Science*. Ed. by L. Ong and R. De Queiroz. Springer, pp. 168–190.

— (2014). "The Logic of Justified Belief, Explicit Knowledge, and Conclusive Evidence". In: *Annals of Pure and Applied Logic* 165.1, pp. 49–81.

Baltag, A. and S. Smets (2011). "Keep Changing Your Beliefs, Aiming for the Truth". In: *Erkenntnis* 75.2, pp. 255–270.

Barrett, J. (2022). "Deviating From the Ideal". In: *Philosophy and Phenomenological Research* 107.1, pp. 31–52.

Belardinelli, G. and B. Schipper (2024). "Implicit Knowledge in Unawareness Structures". In: *Synthese* 204.5, pp. 1–41.

Benavoli, A. and A. Antonucci (2009). "Aggregating imprecise probabilistic knowledge". In: *Proceedings of the 6th International Symposium on Imprecise Probability: Theories and Applications*, n/a–n/a.

Bendaña, J. and E. Mandelbaum (2021). "The Fragmentation of Belief". In: *The Fragmented Mind*. Ed. by C. Borgoni, D. Kindermann, and A. Onofri. Oxford University Press, pp. 78–107.

Bergmann, M. (2018). "Externalist Versions of Evidentialism". In: *Believing in Accordance with the Evidence: New Essays on Evidentialism*. Ed. by McCain K. Springer, pp. 109–123.

Bernoulli, D. (1738/1954). "Exposition of a New Theory on the Measurement of Risk". In: *Econometrica* 22.1, pp. 23–36.

Bjerring, J. (2013). "Impossible Worlds and Logical Omniscience: An Impossibility Result". In: *Synthese* 190.13, pp. 2505–2524.

Bjorndahl, A. and A. Özgün (2020). "Logic and Topology for Knowledge, Knowability, and Belief". In: *Review of Symbolic Logic* 13.4, pp. 748–775.

Bokulich, A. (2011). "How Scientific Models Can Explain". In: *Synthese* 180.1, pp. 33–45.

Bonanno, G., J. Delgrande, and H. Rott (2012). "Special Issue on Multiple Belief Change, Iterated Belief Change and Preference Aggregation". In: *Journal of Philosophical Logic* 41.1.

Bonjour, L. (1980). "Externalist Theories of Empirical Knowledge". In: *Midwest Studies in Philosophy* 5.1, pp. 53–74.

Borgoni, C. (2021). "Rationality in Fragmented Belief Systems". In: *The Fragmented Mind*. Ed. by C. Borgoni, D. Kindermann, and A. Onofri. Oxford University Press, pp. 137–155.

Borgoni, C., D. Kindermann, and A. Onofri (2021). *The Fragmented Mind*. Ed. by C. Borgoni, D. Kindermann, and A. Onofri. Oxford University Press.

Bouton, M. (2004). "Context and behavioral processes in extinction". In: *Learning Memory* 11.5, pp. 485–94.

Bovens, L. and S. Hartmann (2003). *Bayesian Epistemology*. Ed. by S. Hartmann. Oxford University Press.

Bradley, R. (2009). "Revising Incomplete Attitudes". In: *Synthese* 171.2, pp. 235–256.

Bradley, R., C. Helgeson, and B. Hill (2017). "Climate Change Assessments: Confidence, Probability, and Decision". In: *Philosophy of Science* 84.3, pp. 500–522.

Bradley, S. (2019). "Imprecise Probabilities". In: *The Stanford Encyclopedia of Philosophy*. Ed. by E. N. Zalta. Spring 2019. Metaphysics Research Lab, Stanford University.

Bratman, M. (1999). *Faces of Intention: Selected Essays on Intention and Agency*. Cambridge University Press.

Bustamante, J., M. Uengoer, and H. Lachnit (2016). "Reminder Cues Modulate the Renewal Effect in Human Predictive Learning". In: *Frontiers in Psychology* 7, n/a–n/a.

Carnap, R. (1947). "On the Application of Inductive Logic". In: *Philosophy and Phenomenological Research* 8.1, pp. 133–148.

Cevolani, G. and G. Schurz (2017). "Probability, Approximate Truth, and Truthlikeness: More Ways Out of the Preface Paradox". In: *Australasian Journal of Philosophy* 95.2, pp. 209–225.

Cherniak, C. (1983). "Rationality and the Structure of Human Memory". In: *Synthese* 57.2, pp. 163–186.

— (1986). *Minimal Rationality*. MIT Press.

Ciardelli, I., J. Groenendijk, and F. Roelofsen (2018). *Inquisitive Semantics*. Ed. by J. Groenendijk and F. Roelofsen. Oxford University Press.

Clarke, R. (2017). "Preface Writers are Consistent". In: *Pacific Philosophical Quarterly* 98.3, pp. 363–381.

Cohen, L. (1989). "Belief and Acceptance". In: *Mind* 98.391, pp. 367–389.

— (1992). *An Essay on Belief and Acceptance*. Clarendon Press.

Cohen, S. (1999). "Contextualism, Skepticism, and the Structure of Reasons". In: *Philosophical Perspectives* 13, pp. 57–89.

Comesaña, J. (2010). "Evidentialist Reliabilism". In: *Noûs* 44.4, pp. 571–600.

Conee, E. (2004). *Evidentialism: Essays in Epistemology*. Oxford University Press.

Cottingham, J., R. Stoothoff, and D. Murdoch (1984). *The Philosophical Writings of Descartes: Volume 2*. Cambridge University Press.

Cresto, E. (2010). "Belief and Contextual Acceptance". In: *Synthese* 177.1, pp. 41–66.

Crupi, V. (2021). "Confirmation". In: *The Stanford Encyclopedia of Philosophy*. Ed. by E. N. Zalta. Spring 2021. Metaphysics Research Lab, Stanford University.

Davidson, D. (1973). "Radical Interpretation". In: *Dialectica* 27.1, pp. 313–328.

— (1982a). "Paradoxes of Irrationality". In: *Problems of Rationality*. Oxford University Press, pp. 169–188.

— (1982b). *Problems of Rationality*. Oxford University Press.

De Finetti, B. (1970). *Theory of Probability: A Critical Introductory Treatment*. John Wiley.

De Houwer, J. (2009). "The Propositional Approach to Associative Learning as an Alternative for Association Formation models". In: *Learning behavior* 37.1, pp. 1–20.

DeGroot, M. (1974). "Reaching a Consensus". In: *Journal of the American Statistical Association* 69.345, pp. 118–121.

Devine, P. et al. (2012). "Long-term reduction in implicit race bias: A prejudice habit-breaking intervention". In: *Journal of experimental social psychology* 48.6, pp. 1267–1278.

Dietrich, F. and C. List (2016). "Probabilistic Opinion Pooling". In: *Oxford Handbook of Philosophy and Probability*. Ed. by A. Hajek and C. Hitchcock. Oxford University Press, pp. 519–542.

Dietrich, F. and K. Spiekermann (2024). "Deliberation and the Wisdom of Crowds". In: *Economic Theory* 79.2, pp. 603–655.

Dogramaci, S. (2018). "Solving the Problem of Logical Omniscience". In: *Philosophical Issues* 28.1, pp. 107–128.

Dorst, K. (2019). "Evidence: A Guide for the Uncertain". In: *Philosophy and Phenomenological Research* 100.3, pp. 586–632.

— (2022). "Be Modest: You're Living on the Edge". In: *Analysis* 81.4, pp. 611–621.

Dorst, K. and M. Mandelkern (2021). "Good Guesses". In: *Philosophy and Phenomenological Research* 105.3, pp. 581–618.

Douglas, H. (2000). "Inductive Risk and Values in Science". In: *Philosophy of Science* 67.4, pp. 559–579.

Douven, I. (2002). "A New Solution to the Paradoxes of Rational Acceptability". In: *The British Journal for the Philosophy of Science* 53.3, pp. 391–410.

— (2020). "The Ecological Rationality of Explanatory Reasoning". In: *Studies in History and Philosophy of Science Part A* 79.C, pp. 1–14.

Dunn, J. (2019). "Accuracy, Verisimilitude, and Scoring Rules". In: *Australasian Journal of Philosophy* 97.1, pp. 151–166.

Earman, J. (1992). *Bayes or Bust?: A Critical Examination of Bayesian Confirmation Theory*. MIT Press.

Egan, A. (2008). "Seeing and Believing: Perception, Belief Formation and the Divided Mind". In: *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition* 140.1, pp. 47–63.

Egan, A. and M. Davies (2013). "Delusion: Cognitive Approaches–Bayesian Inference and Compartmentalisation". In: *The Oxford handbook of philosophy and psychiatry*. Ed. by K. W. M. Fulford et al. Oxford University Press, pp. 689–727.

Elga, A. and A. Rayo (2021). "Fragmentation and Information Access". In: *The Fragmented Mind*. Ed. by C. Borgoni, D. Kindermann, and A. Onofri. Oxford University Press, pp. 37–53.

— (2022). "Fragmentation and Logical Omniscience". In: *Noûs* 56.3, pp. 716–741.

Ellsberg, D. (1961). "Risk, Ambiguity, and the Savage Axioms". In: *The Quarterly Journal of Economics* 75.4, pp. 643–669.

Evans, V. and M. Green (2006). *Cognitive Linguistics: An Introduction*. Edinburgh University Press.

Fantl, J. and M. McGrath (2009). *Knowledge in an Uncertain World*. Ed. by M. McGrath. Oxford University Press.

Field, C. and B. Jacinto (2022). "Bridge Principles and Epistemic Norms". In: *Erkenntnis* 89.4, pp. 1629–1681.

Fillmore, C. (2006). “Frame semantics and the nature of language”. In: *Annals of the New York Academy of Sciences* 280.1, pp. 20–32.

Fitelson, B. (1999). “The Plurality of Bayesian Measures of Confirmation and the Problem of Measure Sensitivity”. In: *Philosophy of Science* 66.3, pp. 362–378.

— (2021). “A Problem for Confirmation Measure Z”. In: *Philosophy of Science* 88.4, pp. 726–730.

Fitelson, B. and K. Easwaran (2015). “Accuracy, Coherence and Evidence”. In: *Oxford Studies in Epistemology*. Vol. 5. Oxford University Press, pp. 61–96.

Fleisher, W. (2023). “Fragmentation and Old Evidence”. In: *Episteme* 20.3, pp. 542–567.

Fodor, J. (1983). *The Modularity of Mind: An Essay on Faculty Psychology*. MIT Press.

Foley, R. (1979). “Justified Inconsistent Beliefs”. In: *American Philosophical Quarterly* 16.4, pp. 47–257.

— (2009). “Beliefs, Degrees of Belief, and the Lockean Thesis”. In: *Degrees of Belief*. Ed. by F. Huber and C. Schmidt-Petri. Springer, pp. 37–47.

Friedman, M. and L. Savage (1948). “The Utility Analysis of Choices Involving Risk”. In: *Journal of Political Economy* 56.4, pp. 263–292.

Frigg, R. and S. Hartmann (2024). “Models in Science”. In: *The Stanford Encyclopedia of Philosophy*. Ed. by E. N. Zalta and U. Nodelman. Fall 2024. Metaphysics Research Lab, Stanford University.

Frigg, R. and J. Nguyen (2019). “Mirrors Without Warnings”. In: *Synthese* 198.3, pp. 2427–2447.

— (2021). “Scientific Representation”. In: *The Stanford Encyclopedia of Philosophy*. Ed. by E. N. Zalta and U. Nodelman. Winter 2021. Metaphysics Research Lab, Stanford University.

Fumerton, R. (2002). “Theories of Justification”. In: *The Oxford Handbook of Epistemology*. Ed. by P. K. Moser. Oxford University Press, pp. 204–233.

Gabriel, N. and C. O’Connor (2024). “Can Confirmation Bias Improve Group Learning?” In: *Philosophy of Science* 91.2, pp. 1–21.

Galton, F. (1907). "Vox Populi (the Wisdom of Crowds)". In: *Nature* 75.7, pp. 450–451.

Gärdenfors, P. (1988). *Knowledge in Flux. Modelling the Dymanics of Epistemic States*. MIT Press.

Genest, C. and J. Zidek (1986). "Combining Probability Distributions: A Critique and Annotated Bibliography". In: *Statistical Science* 1.1, pp. 114–135.

Gerken, M. (2018). "Against Knowledge-First Epistemology". In: *Knowledge-First Approaches in Epistemology and Mind*. Ed. by E., B. A. Gordon, and Jarvis Carter. Oxford University Press, pp. 46–71.

Gettier, E. (1963). "Is Justified True Belief Knowledge?" In: *Analysis* 23.6, pp. 121–123.

Gigerenzer, G. (2006). "Bounded and Rational". In: *Contemporary Debates in Cognitive Science*. Ed. by R. Stainton. Wiley-Blackwell, pp. 115–133.

Gilbert, D., R. Tafarodi, and P. Malone (1993). "You Can't Not Believe Everything You Read". In: *Journal of Personality and Social Psychology - PSP* 65.2, pp. 221–233.

Gillies, D. (1986). "In Defense of the Popper-Miller Argument". In: *Philosophy of Science* 53.1, pp. 110–113.

Goldberg, S. (2018). *To the Best of Our Knowledge: Social Expectations and Epistemic Normativity*. Oxford University Press.

Goldman, A. (1979). "What is Justified Belief?" In: *Justification and Knowledge*. Ed. by G. Pappas. D. Reidel, pp. 1–23.

— (2005). "Systems-Oriented Social Epistemology". In: *Oxford Studies in Epistemology*. Ed. by T. Szabó-Gendler and J. Hawthorne. Oxford University Press, pp. 189–214.

— (2011). "Toward a Synthesis of Reliabilism and Evidentialism? Or: Evidentialism's Troubles, Reliabilism's Rescue Package". In: *Evidentialism and its Discontents*. Ed. by T. Dougherty. Oxford University Press, pp. 254–280.

— (2012). *Reliabilism and Contemporary Epistemology: Essays*. Oxford University Press.

Good, I. (1966). "On the Principle of Total Evidence". In: *British Journal for the Philosophy of Science* 17.4, pp. 319–321.

— (1983). *Good Thinking: The Foundations of Probability and its Applications*. University of Minnesota Press.

Greco, D. (2015). "Iteration and Fragmentation". In: *Philosophy and Phenomenological Research* 91.3, pp. 656–673.

— (2023). *Idealization in Epistemology: A Modest Modeling Approach*. Oxford University Press.

Grove, A. (1988). "Two Modellings for Theory Change". In: *Journal of Philosophical Logic* 17.2, pp. 157–170.

Harman, G. (1986). *Change in View: Principles of Reasoning*. MIT Press.

Hartmann, S. and J. Sprenger (2010). "Bayesian Epistemology". In: *The Routledge Companion to Epistemology*. Ed. by S. Bernecker and D. Pritchard. Routledge, pp. 609–620.

Hawke, P., A. Özgün, and F. Berto (2020). "The Fundamental Problem of Logical Omniscience". In: *Journal of Philosophical Logic* 49.4, pp. 727–766.

Hawthorne, J. P. (2003). *Knowledge and Lotteries*. Oxford University Press.

Hempel, C. (1965). *Aspects of Scientific Explanation and Other Essays in the Philosophy of Science*. The Free Press.

— (1966). *Philosophy of Natural Science*. Prentice-Hall.

Hintikka, K. (1962). *Knowledge and Belief: An Introduction to the Logic of the Two Notions*. Cornell University Press.

Hintikka, K. J. (1969). "Semantics for Propositional Attitudes". In: *Models for Modalities. Synthese Library*. Vol. 23. Springer, pp. 87–111.

Hoffmann, A. (2012). "Are Propositions Sets of Possible Worlds?" In: *Analysis* 72.3, pp. 449–455.

Icard, T. (2018). "Bayes, Bounds, and Rational Analysis". In: *Philosophy of Science* 85.1, pp. 79–101.

IPCC, Core Writing Team (2023). "Climate Change 2023: Synthesis Report. Contribution of Working Groups I, II and III to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change". In: *IPCC Assessment Report 2023*. Ed. by H. Lee and J. Romero, pp. 1–34.

Jackson, E. and M. Turnbull (2024). "Permissivism, Underdetermination, and Evidence". In: *The Routledge Handbook of the Philosophy of Evidence*. Ed. by C. Littlejohn and M. Lasonen-Aarnio. Routledge, pp. 358–370.

Jeffrey, R. (1956). "Valuation and Acceptance of Scientific Hypotheses". In: *Philosophy of Science* 23.3, pp. 237–246.

— (1983). "Bayesianism with a Human Face". In: *Testing Scientific Theories*. Ed. by J. Earman. Minneapolis: University of Minnesota Press, pp. 133–156.

Johnson, H. and C. Seifert (1994). "Sources of the Continued Influence Effect: When Misinformation in Memory Affects Later Inferences". In: *Journal of Experimental Psychology: Learning, Memory, and Cognition* 20.6, pp. 1420–1436.

Joyce, J. (1998). "A Nonpragmatic Vindication of Probabilism". In: *Philosophy of Science* 65.4, pp. 575–603.

— (2005). "How Degrees of Belief Reflect Evidence". In: *Philosophical Perspectives* 19.1, pp. 153–179.

— (2010). "A Defense of Imprecise Credences in Inference and Decision Making1". In: *Philosophical Perspectives* 24.1, pp. 281–323.

Kahneman, D. and A. Tversky (1979). "Prospect Theory: An Analysis of Decision under Risk". In: *Econometrica* 97.2, pp. 263–292.

Kaplan, M. (1983). "Decision Theory as Philosophy". In: *Philosophy of Science* 50.4, pp. 549–577.

Kauss, D. (2021). "Context-Sensitivity and the Preface Paradox for Credence". In: *Synthese* 199.3–4, pp. 7303–7320.

Kelly, T. (2016). "Evidence". In: *The Stanford Encyclopedia of Philosophy*. Ed. by E. N. Zalta. Winter 2016. Metaphysics Research Lab, Stanford University.

Kemeny, J. (1955). "Fair Bets and Inductive Probabilities". In: *Journal of Symbolic Logic* 20.3, pp. 263–273.

Keynes, J. (1921). *A Treatise on Probability*. Macmillan & co.

Kindermann, D. and A. Onofri (2021). "The Fragmented Mind: An Introduction". In: *The Fragmented Mind*. Ed. by C. Borgoni, D. Kindermann, and A. Onofri. Oxford University Press, pp. 1–36.

King, J. (2024). "Structured Propositions". In: *The Stanford Encyclopedia of Philosophy*. Ed. by E. N. Zalta and U. Nodelman. Winter 2024. Metaphysics Research Lab, Stanford University.

Kolodny, N. (2008). "Why Be Disposed to Be Coherent?" In: *Ethics* 118.3, pp. 437–463.

Konek, J. (2022). "Degrees of Incoherence, Dutch Bookability & Guidance Value". In: *Philosophical Studies* 180.2, pp. 395–428.

Koopman, B. (1940). "The Axioms and Algebra of Intuitive Probability". In: *Annals of Mathematics* 41.2, pp. 269–292.

Kraus, S., D. Lehmann, and M. Magidor (1990). "Nonmonotonic Reasoning, Preferential Models and Cumulative Logics". In: *Artificial Intelligence* 44.1-2, pp. 167–207.

Kyburg, H. (1968). "Bets and Beliefs". In: *American Philosophical Quarterly* 5.1, pp. 54–63.

— (1983). *Epistemology and Inference*. University of Minnesota Press.

Kyburg, H. and M. Pittarelli (1996). "Set-Based Bayesianism". In: *Ieee Transactions on Systems, Man and Cybernetics A* 26.3, pp. 324–339.

Kyburg, H. and C. Teng (1999). "Choosing Among Interpretations of Probability". In: *Proceedings of the Fifteenth Conference on Uncertainty in Artificial Intelligence*. Morgan Kaufmann, pp. 359–365.

— (2001). *Uncertain Inference*. Cambridge University Press.

Lackey, J. (2016). "What Is Justified Group Belief?" In: *Philosophical Review Recent Issues* 125.3, pp. 341–396.

Ladyman, J., J. Lambert, and K. Wiesner (2013). "What is a Complex System?" In: *European Journal for Philosophy of Science* 3.1, pp. 33–67.

Lai, C. et al. (2016). "Reducing Implicit Racial Preferences: II Intervention Effectiveness Across Time". In: *Journal of Experimental Psychology: General* 145.8, pp. 1001–1016.

Lakoff, G. and E. Wehling (2012). *The Little Blue Book. The Essential Guide to Thinking and Talking Democratic*. Free Press.

Lehmann, D. and M. Magidor (1992). "What Does a Conditional Knowledge Base Entail?" In: *Artificial Intelligence* 55.1, pp. 1–60.

Leitgeb, H. (2017). *The Stability of Belief: How Rational Belief Coheres with Probability*. Oxford University Press.

Leitgeb, H. and K. Segerberg (2007). "Dynamic Doxastic Logic: Why, How, and Where To?" In: *Synthese* 155.2, pp. 167–190.

Levi, I. (1974). "On Indeterminate Probabilities". In: *Journal of Philosophy* 71.13, pp. 391–418.

— (1980). *The Enterprise of Knowledge: An Essay on Knowledge, Credal Probability, and Chance*. MIT Press.

— (1985). "Imprecision and Indeterminacy in Probability Judgment". In: *Philosophy of Science* 52.3, pp. 390–409.

— (1986). *Hard Choices: Decision Making Under Unresolved Conflict*. Cambridge University Press.

Levy, N. and E. Mandelbaum (2014). "The Powers That Bind : Doxastic Voluntarism and Epistemic Obligation". In: *The Ethics of Belief: Individual and Social*. Ed. by R. Vitz and J. Matheson. Oxford University Press, pp. 12–33.

Lewis, D. (1973). *Counterfactuals*. Blackwell.

— (1982). "Logic for Equivocators". In: *Noûs* 16.3, pp. 431–441.

— (1980 / 2010). "A Subjectivist?s Guide to Objective Chance". In: *Philosophy of Probability: Contemporary Readings*. Ed. by A. Eagle. Routledge, pp. 263–293.

Lin, H. (2024). "Bayesian Epistemology". In: *The Stanford Encyclopedia of Philosophy*. Ed. by E. N. Zalta and U. Nodelman. Summer 2024. Metaphysics Research Lab, Stanford University.

List, C. (2012). "Judgment Aggregation: A Short Introduction". In: *Handbook of the Philosophy of Economics*. Ed. by U. Mäki. Elsevier, n/a = n/a.

— (2017). "Levels: Descriptive, Explanatory, and Ontological". In: *Noûs* 53.4, pp. 852–883.

— (2022). “Social Choice Theory”. In: *The Stanford Encyclopedia of Philosophy*. Ed. by E. N. Zalta and U. Nodelman. Winter 2022. Metaphysics Research Lab, Stanford University.

List, C. and M. Pivato (2015). “Emergent Chance”. In: *Philosophical Review* 124.1, pp. 119–152.

List, C. and B. Polak (2010). “Introduction to Judgment Aggregation”. In: *Journal of Economic Theory* 145.2, pp. 441–466.

Luce, R. and H. Raiffa (1958). “Games and Decisions: Introduction and Critical Survey”. In: *Philosophy and Phenomenological Research* 19.1, pp. 122–123.

Mahtani, A. (2019). “Imprecise Probabilities”. In: *The Open Handbook of Formal Epistemology*. Ed. by R. Pettigrew and J. Weisberg. PhilPapers Foundation, pp. 107–130.

Makinson, D. (1965). “The Paradox of the Preface”. In: *Analysis* 25.6, pp. 205–207.

Mandelbaum, E. (2014). “Thinking is Believing”. In: *Inquiry: An Interdisciplinary Journal of Philosophy* 57.1, pp. 55–96.

— (2015). “Attitude, Inference, Association: On the Propositional Structure of Implicit Bias”. In: *Noûs* 50.3, pp. 629–658.

Mandelbaum, E. and D. Ripley (2012). “Explaining the Abstract/Concrete Paradoxes in Moral Psychology: The Nbar Hypothesis”. In: *Review of Philosophy and Psychology* 3.3, pp. 351–368.

Mayo-Wilson, C., K. Zollman, and D. Danks (2013). “Wisdom of the Crowds Vs. Groupthink: Learning in Groups and in Isolation”. In: *International Journal of Game Theory* 42.3, pp. 695–723.

Meacham, C. (2016). “Ur-Priors, Conditionalization, and Ur-Prior Conditionalization”. In: *Ergo: An Open Access Journal of Philosophy* 3.17.

Menzel, C. (2024). “Possible Worlds”. In: *The Stanford Encyclopedia of Philosophy*. Ed. by E. N. Zalta and U. Nodelman. Summer 2024. Metaphysics Research Lab, Stanford University.

Miller, E. (2019). “Liars, Tigers, and Bearers of Bad News, Oh My!: Towards a Reasons Account of Defeat”. In: *Philosophical Quarterly* 69.274, pp. 82–99.

Moser, P. and J. Tlumak (1985). "Two Paradoxes of Rational Acceptance". In: *Erkenntnis* 23.2, pp. 127–141.

Nau, R. (2002). "The Aggregation of Imprecise Probabilities". In: *Journal of Statistical Planning and Inference* 105.1, pp. 265–282.

Nozick, R. (1994). *The Nature of Rationality*. Princeton University Press.

O'Connor, C., S. Goldberg, and A. Goldman (2024). "Social Epistemology". In: *The Stanford Encyclopedia of Philosophy*. Ed. by E. N. Zalta and U. Nodelman. Summer 2024. Metaphysics Research Lab, Stanford University.

Okasha, S. and K. Thébault (2020). "Is There a Bayesian Justification of Hypothetico-Deductive Inference?" In: *Noûs* 54.4, pp. 774–794.

Pacuit, E. (2019). "Voting Methods". In: *The Stanford Encyclopedia of Philosophy*. Ed. by E. N. Zalta. Fall 2019. Metaphysics Research Lab, Stanford University.

Pagin, P. and N. Marsili (2021). "Assertion". In: *The Stanford Encyclopedia of Philosophy*. Ed. by E. N. Zalta. Winter 2021. Metaphysics Research Lab, Stanford University.

Parikh, R. (1987). "Knowledge and the Problem of Logical Omniscience". In: *Proceedings of the Second International Symposium on Methodologies for intelligent systems*, pp. 432–439.

Parker, W. (2014). "Values and Uncertainties in Climate Prediction, Revisited". In: *Studies in History and Philosophy of Science Part A* 46, pp. 24–30.

Pettigrew, R. (2016). *Accuracy and the Laws of Credence*. Oxford University Press.

— (2021). "What is Justified Credence?" In: *Episteme* 18.1, pp. 16–30.

Pillin, H. (2025). "Fragmented justification". In: *Synthese*, n/a–n/a.

Popper, K. (1934). *The Logic of Scientific Discovery*. Ed. by Hutchinson Publishing Group. Routledge.

Popper, K. and D. Miller (1983). "A proof of the impossibility of inductive probability". In: *Nature* 302, pp. 687–688.

Porot, N. and E. Mandelbaum (2020). "The science of belief: A progress report". In: *Wiley Interdisciplinary Reviews: Cognitive Science* 12.2, n/a–n/a.

Quine, W. V. O. (1951). "Two Dogmas of Empiricism". In: *Philosophical Review* 60.1, pp. 20–43.

Quintana, I. (2024). "Radical Pooling and Imprecise Probabilities". In: *Erkenntnis* 3, pp. 1–28.

Ramsey, F. (1931). *Foundations: Essays in Philosophy, Logic, Mathematics, and Economics*. Ed. by F. Ramsey and D. Mellor. Humanties Press; Routledge.

Rauhut, A., B. Thomas, and J. Ayres (2001). "Treatments that weaken Pavlovian conditioned fear and thwart its renewal in rats: Implications for treating human phobias". In: *Journal of experimental psychology. Animal behavior processes* 27, pp. 99–114.

Rayo, A. (2013). *The Construction of Logical Space*. Oxford, England: Oxford University Press.

Rendsvig, R., J. Symons, and Y. Wang (2024). "Epistemic Logic". In: *The Stanford Encyclopedia of Philosophy*. Ed. by E. N. Zalta and U. Nodelman. Summer 2024. Metaphysics Research Lab, Stanford University.

Rott, H. (2001). *Change, Choice and Inference: A Study of Belief Revision and Nonmonotonic Reasoning*. Oxford University Press.

Rudner, R. (1953). "The Scientist Qua Scientist Makes Value Judgments". In: *Philosophy of Science* 20.1, pp. 1–6.

Rysiew, P. (2021). "Epistemic Contextualism". In: *The Stanford Encyclopedia of Philosophy*. Ed. by E. N. Zalta. Spring 2021. Metaphysics Research Lab, Stanford University.

Savage, L. (1967). "Difficulties in the Theory of Personal Probability". In: *Philosophy of Science* 34.4, pp. 305–310.

Schervish, M., T. Seidenfeld, and J. Kadane (2002). "Measuring Incoherence". In: *Sankhyā: The Indian Journal of Statistics, Series A (1961-2002)* 64.3, pp. 561–587.

Schmidt, E. (2020). "The Explanatory Merits of Reasons-First Epistemology". In: *Concepts in Thought, Action, and Emotion: New Essays*. Ed. by C. Demmerling and D. Schroeder. Routledge, pp. 75–93.

Schoenfield, M. (2020). "Can Imprecise Probabilities Be Practically Motivated? A Challenge to the Desirability of Ambiguity Aversion". In: *Philosophers' Imprint* 20.30, pp. 1–21.

Schroeder, M. (2021). *Reasons First*. Oxford University Press.

Sedley, D. and A. Long (1980). *Plato: Meno and Phaedo*. Cambridge University Press.

Seidenfeld, T., M. Schervish, and J. Kadane (1995). "A Representation of Partially Ordered Preferences". In: *The Annals of Statistics* 23.6, pp. 2168–2217.

Seifert, C. (2002). "The continued influence of misinformation in memory: What makes a correction effective?" In: *Psychology of Learning and Motivation* 41, pp. 265–292.

Šešelja, D. (2023). "Agent-Based Modeling in the Philosophy of Science". In: *The Stanford Encyclopedia of Philosophy*. Ed. by E. N. Zalta and U. Nodelman. Winter 2023. Metaphysics Research Lab, Stanford University.

Simon, H. (1955). "A Behavioral Model of Rational Choice". In: *The Quarterly Journal of Economics* 69.1, pp. 99–118.

— (1957). *Models of Man*. John Wiley.

Sinnott-Armstrong, W. (2008). "Abstract + Concrete = Paradox". In: *Experimental Philosophy*. Ed. by J. Knobe and S. Nichols. Oxford University Press, pp. 209–230.

Skipper, M. and J. Bjerring (2019). "A Dynamic Solution to the Problem of Logical Omniscience". In: *Journal of Philosophical Logic* 48.3, pp. 501–521.

Skovgaard-Olsen, N. (2017). "The Problem of Logical Omniscience, the Preface Paradox, and Doxastic Commitments". In: *Synthese* 194.3, pp. 917–939.

Skyrms, B. (1980). "Higher Order Degrees of Belief". In: *Prospects for Pragmatism: Essays in Memory of F P Ramsey*. Ed. by David Hugh Mellor. Cambridge University Press, pp. 109–137.

Smith, N. (2025). "Interpreting Imprecise Probabilities". In: *Philosophical Quarterly* n/a.n/a.

Smithies, D. (2015). "Ideal Rationality and Logical Omniscience". In: *Synthese* 192.9, pp. 2769–2793.

Sprenger, J. and S. Hartmann (2019). *Bayesian Philosophy of Science*. Oxford University Press.

Staffel, J. (2018). *Unsettled Thoughts: A Theory of Degrees of Rationality*. Oxford University Press.

Stalnaker, R. (1984). *Inquiry*. Cambridge University Press.

— (1991). "The Problem of Logical Omniscience, I". In: *Synthese* 89.3, pp. 425–440.

— (2006). "On Logics of Knowledge and Belief". In: *Philosophical Studies* 128.1, pp. 169–199.

Stanley, J. (2005). *Knowledge and Practical Interests*. Oxford University Press.

Steele, K. (2012). "The Scientist Qua Policy Advisor Makes Value Judgments". In: *Philosophy of Science* 79.5, pp. 893–904.

Steup, M. and R. Neta (2020). "Epistemology". In: *The Stanford Encyclopedia of Philosophy*. Ed. by E. N. Zalta. Fall 2020. Metaphysics Research Lab, Stanford University.

Stewart, R. and L. Elkin (forthcoming). "Diving Into Fair Pools: Algorithmic Fairness, Ensemble Forecasting, and the Wisdom of Crowds". In: *Analysis*.

Stewart, R. and I. Quintana (2018). "Probabilistic Opinion Pooling with Imprecise Probabilities". In: *Journal of Philosophical Logic* 47.1, pp. 17–45.

Stone, M. (1961). "The Opinion Pool". In: *The Annals of Mathematical Statistics* 23.4, pp. 1339–1342.

Strasser, C. and G. A. Antonelli (2024). "Non-monotonic Logic". In: *The Stanford Encyclopedia of Philosophy*. Ed. by E. N. Zalta and U. Nodelman. Winter 2024. Metaphysics Research Lab, Stanford University.

Surowiecki, J. (2005). *The Wisdom of Crowds*. Anchor.

Swinburne, R. (2001). *Epistemic Justification*. Oxford University Press.

Tentori, K. et al. (2007). "Comparison of Confirmation Measures". In: *Cognition* 103.1, pp. 107–119.

Thomason, R. (1986). "The Context-Sensitivity of Belief and Desire". In: *Reasoning about actions and plans*. Ed. by M. Georgeff and A. Lanksy. Morgan Kaufmann Publishers, pp. 341–360.

— (2001). "Three Interactions Between Context and Epistemic Locutions". In: *Modeling and Using Context*. Ed. by R. Young et al. Springer, pp. 467–481.

Thorstad, D. (2024). "Why Bounded Rationality (in Epistemology)?" In: *Philosophy and Phenomenological Research* 108.2, pp. 396–413.

Turri, J. (2010). "On the Relationship Between Propositional and Doxastic Justification". In: *Philosophy and Phenomenological Research* 80.2, pp. 312–326.

Vallinder, A. (2018). "Imprecise Bayesianism and Global Belief Inertia". In: *British Journal for the Philosophy of Science* 69.4, pp. 1205–1230.

Van Benthem, J. (1996). *Exploring logical dynamics*. CSLI Publications.

— (2006). "Epistemic Logic and Epistemology: The State of Their Affairs". In: *Philosophical Studies* 128.1, pp. 49–76.

— (2007). "Dynamic Logic for Belief Revision". In: *Journal of Applied Non-Classical Logics* 17.2, pp. 129–155.

Van Benthem, J. and E. Pacuit (2011). "Dynamic Logics of Evidence-Based Beliefs". In: *Studia Logica* 99.1-3, pp. 61–92.

Van Ditmarsch, H., W. Van der Hoek, and B. Kooi (2007). *Dynamic Epistemic Logic*. Springer.

Van Fraassen, B. (1980). *The Scientific Image*. Oxford University Press.

Von Neumann, J. and O. Morgenstern (1944). *Theory of Games and Economic Behavior*. Princeton University Press.

Von Wright, G. (1951). *An Essay in Modal Logic*. North-Holland Publishing Company.

Vul, E and H. Pashler (2008). "Measuring the Crowd Within: Probabilistic Representations Within Individuals". In: *Psychological science* 19.7, pp. 645–647.

Walley, P. (1991). *Statistical Reasoning with Imprecise Probabilities*. Chapman & Hall.

Wedgwood, R. (2008). "Contextualism About Justified Belief". In: *Philosophers' Imprint* 8.9, pp. 213–219.

Wegner, D., G. Coulton, and R. Wenzlaff (1985). "The Transparency of Denial. Briefing in the Debriefing Paradigm". In: *Journal of Personality and Social Psychology* 49.2, pp. 338–346.

Wheeler, G. (2024). "Bounded Rationality". In: *The Stanford Encyclopedia of Philosophy*. Ed. by E. N. Zalta and U. Nodelman. Winter 2024. Metaphysics Research Lab, Stanford University.

White, R. (2005). "Epistemic Permissiveness". In: *Philosophical Perspectives* 19.1, pp. 445–459.

Williams, J. (2021). "Once you think you're wrong, you must be right: new versions of the preface paradox". In: *Synthese* 198.7, pp. 1801–1825.

Williamson, J. (2010). *In Defence of Objective Bayesianism*. Oxford University Press.

Williamson, T. (2000). *Knowledge and its Limits*. Oxford University Press.

Yalcin, S. (2018). "Belief as Question-Sensitive". In: *Philosophy and Phenomenological Research* 97.1, pp. 23–47.

— (2021). "Fragmented but Rational". In: *The Fragmented Mind*. Ed. by C. Borgoni, D. Kindermann, and A. Onofri. Oxford University Press, pp. 156–180.

Zaffalon, M. and G. De Cooman (2005). "Editorial: Imprecise probability perspectives on artificial intelligence". In: *Annals of Mathematics and Artificial Intelligence* 45.1–2, pp. 1–4.